Working Paper

ZACHARY BURDETTE, DWIGHT PHILLIPS, JACOB L. HEIM, EDWARD GEIST, DAVID R. FRELINGER, CHAD HEITZENRATER, KARL P. MUELLER

An AI Revolution in Military Affairs?

How Artificial Intelligence Could Reshape Future Warfare

RAND Global and Emerging Risks

WR-A4004-1
June 2025

RAND working papers are intended to share early insights and solicit informal peer review. This working paper has been approved for circulation by RAND but has not been peer reviewed or professionally edited or proofread. Feedback to the project team is welcome via email. This working paper can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. This working paper does not necessarily reflect the opinions of RAND's research clients and sponsors. RAND® is a registered trademark. Learn more at www.rand.org.

For more information on this publication, visit **www.rand.org/t/WRA4004-1**.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.

© 2025 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, visit www.rand.org/about/publishing/permissions.

WR-A4004-1

About This Working Paper

How will advances in artificial intelligence (AI) shape the future of war? This working paper offers a conceptual framework and preliminary assessment of AI's potential to disrupt how militaries fight and win wars. The framework evaluates how AI could impact four "building block" competitions in military affairs. This building block approach breaks down warfare into component competitions revolving around (1) quantity versus quality, (2) hiding versus finding, (3) centralized versus decentralized command and control, and (4) cyber offense versus cyber defense. The findings suggest that the U.S. military may need to change important aspects of how it traditionally approaches force planning and operational concepts to exploit AI's potential.

Intended Audience

This working paper is intended for policymakers, defense analysts, and technologists considering the potential impact of AI on future military conflicts.

RAND working papers are intended to share early insights and solicit informal peer review. This working paper has been approved for circulation by RAND but has not been peer reviewed or professionally edited or proofread. Feedback to the project team is welcome via email. This working paper can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. This working paper does not necessarily reflect the opinions of RAND's research clients and sponsors.

Technology and Security Policy Center

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was undertaken by the division's Technology and Security Policy Center, which explores how high-consequence, dual-use technologies change the global competition and threat environment, then develops policy and technology options to advance the security of the United States, its allies and partners, and the world. For more information, contact tasp@rand.org.

Funding

This effort was independently initiated and conducted within the Technology and Security Policy Center using income from operations and gifts from philanthropic supporters, which have been made or recommended by DALHAP Investments Ltd., Effektiv Spenden, Ergo Impact, Founders Pledge, Charlottes och Fredriks Stiftelse, Good Ventures, Jaan Tallinn, Longview,

Open Philanthropy, and Waking Up Foundation. A complete list of donors and funders is available at www.rand.org/TASP. RAND donors and grantors have no influence over research findings or recommendations.

Acknowledgments

For helpful feedback, we thank the other members of the Geopolitics of Artificial General Intelligence (AGI) research team, especially Joel Predd, Jim Mitre, Barry Wilson, Casey Mahoney, Tobias Sytsma, and Scott Savitz. We also thank Ryan Bain, Matt Barry, Brad Carson, and Matty Goldman for providing comments on earlier drafts. Sean Atkins provided helpful insights on AI's implications for cyber operations. Hiwot Demelash provided excellent research support. Ryan Bain assisted with formatting the document.

Summary

Issue

How will advances in artificial intelligence (AI) shape the future of war? Despite significant public interest in AI's potential to revolutionize warfare, researchers are still in the early stages of exploring how AI will change warfighting. This report offers a conceptual framework and preliminary assessment to help set the terms for a more systematic debate about AI's potential to change how militaries fight and win wars.

Approach

In our analysis, we assume that AI delivers on its goal of removing many of the limits of human intelligence as a constraint on military operations and examine the consequences of this assumption. The report's framework evaluates how AI could impact four "building block" competitions of military operations. This building block approach identifies component competitions in warfare revolving around (1) quantity versus quality, (2) hiding versus finding, (3) centralized versus decentralized command and control (C2), and (4) cyber offense versus cyber defense. For each building block, we identify ways in which AI could affect key drivers and tradeoffs that will shape the competition.

Key Findings

Based on the analysis of the four building block competitions, we assess that the U.S. military may need to change important aspects of how it traditionally approaches operational concepts and force planning to exploit AI's potential. Since AI technology is still maturing, these are tentative hypotheses:

- **Quantity could gain a significant edge vis-à-vis quality.** As mass becomes cheaper and more capable, the relative cost-effectiveness of investing in quality versus quantity could change. Advances in autonomy and robotics might make it feasible to field large numbers of platforms at a scale that was previously financially impractical without dramatically sacrificing performance. Exquisite platforms and weapons will still have a role in future fights, but the relative advantages of mass in military operations will likely increase.
- **More sophisticated hiding could help offset advances in finding, but this requires new approaches and investments in deception.** Hiding could become much more effective than it is today if militaries use AI to create "fog-of-war machines" that orchestrate sophisticated deception campaigns with large numbers of advanced decoys. However, this potential advantage depends on the kind of information the hider is trying to conceal and the finder's ability to leverage advanced networks of AI-enabled sensors. The result will likely be an especially contested measure-countermeasure race between

hiders and finders in which the relative advantage could vary across contexts. AI and robotic decoys are likely to be particularly useful for hiding great powers' nuclear forces, which reduces the risk that AI will upend strategic stability between peer competitors.

- **Mission command—a hybrid of centralized and decentralized modes—will remain dominant.** The main limiting factor in improving C2 is not better intelligence or faster decision making but instead the resilience of communications links between forces.
- **Cyber defenses will benefit from AI in ways that could make battle networks more resilient against attacks in the long term.** Critically, AI could help address challenges with scale and speed that currently limit the effectiveness of cyber defenses and give the offense an edge. However, the cyber offense will also benefit from AI, and attackers will retain the ability to penetrate networks at certain times and places.
- **Militaries that fail to embrace mass and deception may be at a serious disadvantage.** Relying on small force structures of exquisite capabilities is quickly becoming a liability rather than an asset. Advances in sensing networks and long-range strike mean that militaries must place deception at the center of how they operate. Making this shift will require new approaches to U.S. defense production and sustainment, including protecting the logistics tail required to sustain a much larger force structure over long distances.

Implications

Militaries that fully embrace significant changes to their operational concepts and force structures will have advantages over those that simply use AI to make their existing approaches to warfighting incrementally better. Exploiting AI's military potential will be as much an organizational challenge as a technological one because it will require disruptive changes to how militaries have become comfortable organizing and operating.

- **The United States should begin investing in mass and fog-of-war complexes now.** Opportunities to capitalize on the trends that this working paper identifies are available today even if AI is still maturing. The United States should increase its experiments and investments in key areas such as unmanned systems and advanced decoys. Force planners should also ensure that high-low mixes reflect a greater emphasis on mass than they currently do and that new exquisite systems include plans for deception and stealth from the onset, not as an afterthought.
- **The United States needs to allocate scarce resources under the assumption that it will face sophisticated and adaptive adversaries.** There may be temptations to use AI to double down on existing approaches in the hope that the United States will have an AI first-mover advantage, enabling it to operate in ways that defy the trends we hypothesize. The United States should instead prepare for the case that such advantages could be fleeting and allocate scarce resources toward approaches that will provide durable long-term advantages.
- **The United States needs to ensure the Defense Industrial Base can generate mass.** The Defense Industrial Base currently focuses on building exquisite capabilities in small numbers, and it struggles to do so at scale. To transition to a force structure oriented around mass, the United States will need to reorient and strengthen the industrial base. Doing so will require lead time and resources, so it should start that process now.

Contents

About This Working Paper iii
Summary v
Contents vii
Figures and Tables viii
Chapter 1. Introduction 1
Mechanisms Through Which AI Could Change Warfare 2
A Framework for Analysis: "Building Block" Competitions 6
Chapter 2. Quality vs. Quantity 10
How AI Can Help Both Quantity and Quality in Air Combat 11
How Much Quality is Enough to Overcome Greater Quantity? 12
What are Potential Limiting Factors on Embracing Mass? 17
Chapter 3. Hider vs. Finder 25
How AI Can Help Finders 26
How AI Can Help Hiders 28
When Will AI Help Hiders or Finders More? 33
Chapter 4. Centralized vs. Decentralized C2 37
How AI Can Help Both Centralized and Decentralized C2 Architectures 38
Why Mission Command Will Remain Dominant 39
Chapter 5. Cyber Offense vs. Cyber Defense 46
How AI Can Help Both Attackers and Defenders Find Cyber Vulnerabilities 47
How AI Can Help Detect and Hide an Attacker's Maneuvers and Effects 50
How AI Can Help Patch and Exploit Cyber Vulnerabilities 51
Chapter 6. Does an AI First-Mover Advantage Matter? 53
Chapter 7. Implications: A New AI-Enabled Way of War? 56
Summary of Key Findings 56
Implications for Conventional Warfare 57
Implications for Strategic Stability 61
Implications for Unconventional Warfare 62
Abbreviations 66
References 67
About the Authors 76

Figures and Tables

Figures

Figure 2.1. Quality vs. Quantity Break-Even Points (Using the Lanchester Square Law) 15
Figure 2.2. Quality vs. Quantity Break-Even Points with Historical Example 17
Figure 2.3: Cost Calculations and the Affordability of Quantitative Advantages 19
Figure 2.4: How Autonomy may Change the Cost of Relying on Quantity 20
Figure 2.5: Asymmetries in Distance Increase the Quantity of Aircraft Required 22
Figure 3.1 The Five Links of the Kill Chain 27
Figure 3.2 Illustrative Fog of War Machine 31
Figure 4.1. Situational Awareness vs. Proximity to Tactical Engagement (USAF) 40
Figure 4.2. Notional C2 Decentralization (USAF) 42
Figure 4.3. Inter-flight Communications Increase a Swarm's Ability to Find and Kill Targets 43
Figure 4.4. Inter-flight Communications Increase a Swarm's Robustness to Attrition 44
Figure 5.1. Baseline Cyber Offense vs. Cyber Defense 48
Figure 6.1. First-Mover Advantage Typology 54

Tables

Table 1.1. Ways AI Could Affect Warfare 5
Table 1.2. Hypotheses About the Military Implications of Advanced AI 8
Table 2.1. Lanchester Assumptions: Human Battle vs. Robotic Battle 13

Chapter 1. Introduction

How will advances in artificial intelligence (AI) influence the conduct of war? There is a growing belief among policymakers and analysts that AI will transform the future of war and could even reshape the international balance of power.[1] Despite significant public interest in AI's potential to revolutionize warfare, researchers are still in the early stages of understanding how AI will actually change warfighting.[2] Commentators sometimes posit that AI will lead *deus ex machina* to superintelligence and superweapons. This working paper proposes a conceptual framework to help set the terms for a more systematic debate about AI's implications for the future of war. It breaks down the broad question of how AI will shape warfare into discrete "building blocks" that are central to modern military operations. It then provides a preliminary assessment of how AI could change the dynamics in these critical building block competitions. The central argument is that the characteristics of AI-enabled warfare will likely reward militaries that invest in capabilities and operational concepts that leverage mass, deception, mission command, and network resilience.

Because AI is still an emerging technology, our assessment about how it could shape future military operations is necessarily tentative and rests on assumptions about its future performance. AI performance exists on a spectrum that varies in the quality and breadth of the AI's capability. Today, artificial narrow intelligence (ANI) can handle discrete tasks under specific conditions, sometimes exceeding human performance. In the future, progress toward artificial general intelligence (AGI) could offer the ability to consistently match or exceed human performance across a wide range of cognitive tasks and under many different conditions. This working paper explores the implications for warfare if AI lives up to this vision of *removing the limits of human intelligence as a constraint* on military operations. In many cases, the trends that we explore are already starting to unfold with today's frontier AI combined with robotics, so making this assumption does not imply that investing in new AI-enabled capabilities and concepts is

[1] See, for example, John Tirpak, "Kendall: In US-China 'Race for Technological Superiority,' AI May Be the Key," *Air & Space Forces Magazine*, October 29, 2024; Michèle Flournoy, "AI is Already at War: How Artificial Intelligence Will Transform the Military," *Foreign Affairs*, Vol. 102, No. 6, 2023; Shashank Joshi, "AI Will Transform the Character of Warfare," *The Economist*, June 20, 2024. On the potential implications for the balance of power, see: Michael Horowitz, "Artificial Intelligence, International Competition, and the Balance of Power," *Texas National Security Review*, Vol. 1, No. 3, May 2018, pp. 36–57; and Elliot Ackerman and James Stavridis, "Drone Swarms Are About to Change the Balance of Military Power," *Wall Street Journal*, March 14, 2024.

[2] For examples of earlier work on this topic, see: Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, W.W. Norton & Company, 2018; Paul Scharre, *Four Battlegrounds: Power in the Age of Artificial Intelligence*, New York: W.W. Norton & Company, 2023; Avi Goldfarb and Jon Lindsay, "Prediction and Judgment: Why Artificial Intelligence Increases the Importance of Humans in War," *International Security*, Vol. 46, No. 3, February 25, 2022, pp. 7–50; Benjamin Chang, *Artificial Intelligence and the US-China Balance of Power*, dissertation, Massachusetts Institute of Technology, 2021; and Lance Menthe et al., *Understanding the Limits of Artificial Intelligence for Warfighters: Volume 1, Summary,* RAND Corporation, RR-A1722-1, January 3, 2024.

something that militaries should defer to the future. But making this assumption about AI's potential helps bypass technical debates about the state of the technology today and instead focuses attention on the long-term implications of where the technology might be heading.

Mechanisms Through Which AI Could Change Warfare

Understanding AI's potential to reshape warfare requires identifying specific mechanisms through which the technology could impact how militaries fight or build capabilities. One of AI's defining features is that it is a "general purpose technology" rather than one discrete capability, making it more like electricity or the internal combustion engine than specific platforms or weapons such as tanks or nuclear weapons.[3] Because AI is an enabling technology with many applications, and because different militaries will adopt AI for a wide range of purposes to meet their particular needs, there is no single way that the technology will change warfare.[4] This section therefore identifies a range of ways that militaries might leverage AI.

To identify AI's applications, we distinguish between warfighting and preparing for war.[5] *Warfighting* is the realm of active combat operations, characterized by its dynamic and unpredictable nature, where the fog of war, friction, and complex interactions with the enemy require flexibility and adaptability. AI could change warfighting through several mechanisms:

- **Analyzing information and generating knowledge.** By synthesizing and interpreting vast amounts of data, AI could significantly improve decisionmakers' understanding of the battlespace. Many of the potential cognitive benefits from AI do not depend on AI surpassing human levels of intelligence, just the ability to think about a lot more on a much faster timescale. For example, the U.S. military deployed an AI-enabled image classifier as early as 2017 to sort through huge volumes of intelligence and detect specific objects, such as vehicles, in support of operations against the Islamic State.[6] Leveraging advances in AI to more rapidly and reliably identify specific objects offers advantages including an improved ability to strike high-value targets and better threat warning. There are also potential benefits for logistics and support functions. For example, early tests of AI for predictive maintenance increased mission-capable rates for U.S. military aircraft.[7]
- **Managing complicated actions and systems.** AI could coordinate and synchronize complicate tasks that involve too many entities or components for humans to manage at

[3] Horowitz, 2018, p. 41.

[4] Michael Horowitz, *The Diffusion of Military Power: Causes and Consequences for International Politics*, Princeton, N.J: Princeton University Press, 2010; and Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, Princeton, N.J.: Princeton University Press, 2024.

[5] Similarly, Colin Gray makes a distinction between "War Proper"—i.e., warfighting—and "Preparation for War." Colin S. Gray, *Modern Strategy*, New York: Oxford University Press, 1999. Michael Howard also distinguishes between the operational (war proper) and logistical (preparation for war) dimensions of strategy. Michael Howard, "The Forgotten Dimensions of Strategy," *Foreign Affairs,* Vol. 57, No. 5, Summer 1979.

[6] Scharre, 2023, pp. 56-58.

[7] Scharre, 2023, pp. 196-197.

the necessary speed and scale.[8] For example, researchers are developing AI tools to enable large swarms of autonomous drones to coordinate with each other.[9] Advanced AI might make the challenges of deconflicting crowded airspace and synchronizing the fires of potentially hundreds of drones at once manageable for machines, even if they are too complex for humans to do effectively.

- **Replacing human labor.** AI could make autonomous systems more effective and reliable, potentially replacing human labor across both physical and cognitive tasks. When paired with robotics, this could enable better performance and/or lower risk for dangerous jobs. For example, in 2024, the U.S. military announced that an AI-controlled F-16 put up a "roughly even fight" in an exercise against a crewed F-16 with an expert human pilot.[10] Replacing human labor could be particularly useful for tasks constrained by human physical limitations, including piloting aircraft on missions requiring very long endurance, experiencing high g-loads, and using small airframes. But replacing humans even in mundane tasks could reduce casualties, improve efficiencies, and reduce personnel and training costs.
- **Making better decisions.** AI could provide decision support tools that help military and civilian decisionmakers better assess their strategic options, evaluate different courses of action (COAs), and translate decisions into orders and coordinating guidance.[11] For example, AI is already better than humans at well-structured strategy games such as chess and Go, and human players can reach peak performance when working in concert with AI tools.[12] Future advances in AI could extend the decision-making benefits of human-machine teaming to more ambiguous and dynamic contexts, like war. Advances in AI also have the potential to enable humans to make not only better but also faster decisions, potentially providing an advantage over opponents in contexts where rapid decision making and accelerating tactical adaptation to an adversary's actions are important.[13]

Beyond combat operations, AI could also change how countries prepare for war. *Preparation for war* encompasses activities including force development, technological innovation, materiel production, and logistics. These kinds of support functions are relevant both for building capabilities during peacetime and sustaining forces during protracted wars. AI could have at least two potential impacts on how militaries prepare for wars:

---

[8] Matthew Walsh et al., *"Exploring the Feasibility and Utility of Machine-Learning Assisted Command and Control: Volume 1, Findings and Recommendations*", RAND Corporation, RR-A263-1, 2021.

[9] Scharre, 2018, p. 11.

[10] Greg Hadley, "In F-16 Dogfight, AI and Human Pilots Are 'Roughly an Even Fight,' Says Kendall," *Air and Space Forces Magazine*, May 8, 2024.

[11] Eric Robinson, Daniel Egel, and George Bailey, *"Machine Learning for Operational Decisionmaking in Competition and Conflict: A Demonstration Using the Conflict in Eastern Ukraine*", RAND Corporation, RRA815-1, 2023; David Stebbins et al., *"Exploring Artificial Intelligence Use to Mitigate Potential Human Bias within U.S. Army Intelligence Preparation of the Battlefield Process*", RAND Corporation, RR-A2763-1, 2024.

[12] Scharre, 2018, p. 321.

[13] Bryan Clark, Dan Patt, and Harrison Schramm, *Mosaic Warfare: Exploiting Artificial Intelligence and Autonomous Systems to Implement Decision-Centric Operations*, Center for Strategic and Budgetary Assessments, 2020.

- **Making production more efficient**. AI could optimize supply chains and advance robotic manufacturing, making the production of weapons and platforms faster and cheaper. AI could impact multiple elements of the production chain, including supply chain management, process optimization, asset management, quality control, and automated production.[14] For example, researchers are exploring ways to use AI for end-to-end robotic manufacturing, including streamlining designs and making robotic manufacturing tools more adaptive so they can work with small electronics such as drones and munitions.[15]
- **Designing new military capabilities.** AI could spur technological discovery and advance military research and development (R&D) efforts.[16] As part of an emerging "AI for Science" paradigm, researchers are integrating AI into every stage of their work to speed up the pace of discovery.[17] Scientists have already used AI to achieve Nobel Prize-winning breakthroughs, and AI could assist with a wide range of applied problems in the future, including the design of software, platforms, weapons, and structural materials. Combining AI-enabled design with additive manufacturing like 3D printing could enable more dynamic adjustments to parts, platforms, and munitions—including at the edge of operations.[18]

Table 1.1 provides an overview of these different mechanisms.[19]

[14] Henri Bristol et al., "Generative AI in Manufacturing Lighthouses | McKinsey," McKinsey, April 2024.

[15] Patrick Tucker, "Could this breakthrough in robot manufacturing reshape warfare?," *Defense One*, April 14, 2025.

[16] Gaurav Sett, "How AI Can Automate AI Research and Development," RAND Corporation, October 24, 2024.

[17] Hanchen Wang et al., "Scientific discovery in the age of artificial intelligence," *Nature*, Vol. 620, 2023.

[18] Daniel M. Gerstein and Erin N. Leidy, *"Emerging Technology and Risk Analysis: Additive Manufacturing*", RAND Corporation, RR-A2984-1, February 15, 2024, pp. 3-6.

[19] See, also, Shashank Joshi, "AI on the Battlefield: A Fait Accompli?" *Survival*, Vol. 67, No. 2, 2025, pp. 115-122; Forrest E. Morgan et al., *Military Applications of Artificial Intelligence: Ethical Concerns in an Uncertain World*, RAND Corporation, RR-3139-1-AF, April 28, 2020.

Table 1.1. Ways AI Could Affect Warfare

Realm	Mechanism	Logic	Example Use Cases
Warfighting	Analyzing information and generating knowledge	AI makes it possible to analyze much larger volumes of data and generate information at much faster speeds than is practical for humans	• Data-intensive surveillance • Multi-source sensor fusion • Targeting • Cyber warfare • Electronic warfare
	Managing complicated actions and systems	AI can coordinate and synchronize independent agent activities that are too complex for humans to manage at tactically or operationally relevant timescales	• Managing intelligent robotic swarms • Coordinating airspace for thousands of platforms • Synchronizing kinetic and non-kinetic fires • Managing distributed logistics
	Replacing human labor	AI can make autonomous systems more capable of replacing human labor for cognitive tasks and, when combined with robotics, for physical tasks	• Obstacle belt breaching and clearance • Piloting aircraft and other platforms • Supply convoys • Depot operations
	Making better decisions	AI can make better and faster decisions than humans because of its pattern recognition capabilities and its unconstrained or less biased thinking	• Generating and evaluating courses of action • Managing air and missile defense
Preparing for War	Making production more efficient	AI could optimize supply chains or improve robotic manufacturing, making weapons and materiel production faster and cheaper	• Automated production lines • AI-assisted expertise in intensive labor in areas such as assembly and quality control
	Designing new military capabilities	AI could accelerate R&D, particularly for compute- or cognitive-intensive technological solutions for military applications	• Improved research in capabilities like directed-energy weapons, bio-defense, and advanced structural material performance (e.g. engines, armor)

SOURCE: RAND analysis.
NOTE: These reflect assumptions about AI's potential capabilities based on first principles rather than research findings based on technical analysis or modeling.

Just as important as identifying what AI could do is acknowledging what it could not do. Throughout the analysis, we assume that AI faces fundamental limits rooted in the laws of physics. For example, aircraft have limitations based on size, weight, and power (SWaP) constraints in how many weapons they could plausibly carry, or how far and fast they could fly. AI may accelerate innovation that leads to new technologies, such as more efficient energy storage and usage, but it will not enable militaries to develop capabilities or operate in ways that simply defy the laws of physics.

Finally, AI's potential impact depends in part on how it interacts with other existing and emerging technologies. We consider throughout this working paper how combing AI with other technologies will enable militaries to operate in ways that would not work well with either technology in isolation. For example, by some estimates Ukraine lost as many as 10,000 drones per month to Russian jamming in 2023.[20] Advances in AI raise the possibility that even cheap drones will no longer require constant human control via vulnerable radio-frequency links. Conversely, many AI applications depend on robotics. Without being paired with hardware, AI is primarily useful as an analytical and administrative tool and lacks the ability to maneuver in the physical world or manipulate physical objects.

A Framework for Analysis: "Building Block" Competitions

Some analysts predict that advances in AI will produce a *revolution in military affairs* (RMA), which is a discontinuous jump in military effectiveness after which militaries that embrace new capabilities and operational concepts gain decisive advantages over adversaries that have not yet adapted.[21] What makes these changes revolutionary is that they offer dramatic improvements to operational performance, not just the linear and incremental improvements that results from most innovations. Examples of past RMAs in the 20th century include carrier warfare supplanting battleships for naval combat and doctrinal innovation in mechanized combined arms warfare that enabled blitzkrieg's penetrating breakthroughs.[22] In these cases and in other RMAs, identifying the right concepts for employing the capability is as essential as the technological breakthrough itself.

Instead of litigating whether and when AI will produce revolutionary changes to the conduct of war, our framework focuses on identifying *how* militaries should fight in a new AI-enabled paradigm. If AI is going to cause a larger revolution in warfare, it will rest on the micro-foundations of how AI changes specific ways that militaries operate. This working paper's framework looks at how AI's ability to accomplish the tasks described above will impact "building block" competitions in military affairs. The building block approach identifies component military-technological competitions that are foundational to how modern militaries fight and win wars.[23]

---

[20] Jack Watling, Oleksandr V. Danylyuk, and Nick Reynolds, *Preliminary Lessons from Ukraine's Offensive Operations, 2022–*23, Royal United Services Institute, July 18, 2024.

[21] Owen Daniels, *The "AI RMA" The Revolution Has Not Arrived (Yet)*, Andrew Marshall Foundation, 2022; Williamson Murray and Macgregor Knox, "Thinking about Revolutions in Warfare," in *The Dynamics of Military Revolution:1300-2050,* Cambridge University Press, 2001, 12.

[22] Williamson Murray and Allan Millet, *Military Innovation in the Interwar Period*, Cambridge University Press, 1998.

[23] For another recent example of the building block approach, see: Andrew Krepinevich, *The Origins of Victory: How Disruptive Military Innovation Determines the Fates of Great Powers*, Yale University Press, 2023.

A key choice in using this methodology is selecting which building block competitions to examine. Our criteria prioritized competitions that will remain central to warfare even as technology evolves and that are generalizable across domains. Warfare increasingly revolves around contests of "battle networks," where multiple military units act in concert through networked operations to achieve objectives.[24] Modern battle networks use sensors to find targets, shooters to strike them, and communications links to stitch the network together.[25] The ability to make good and timely decisions based on the information flowing through these networks is also critical to their effectiveness.[26]

We identified four "building block" competitions that represent key overarching themes: quantity versus quality, hiding versus finding, centralized versus decentralized command and control (C2), and the resilience of cyber network defenses against cyber-attacks. These competitions reflect central questions about what battle networks will look like and how they will operate in the future. How will they see, understand, and strike targets? What is the mix of attributes and the quantities of sensors and shooters that will make the networks more effective? How should militaries direct and synchronize the actions of thousands of nodes in future battle networks? How will they protect themselves from the adversary's battle networks, including both cyber-attacks and long-range precision strike?

The framework focuses on identifying relative advantages within these building block competitions. There is a tendency to focus on how AI will make everything more capable in absolute terms, but wars are relative contests. Analysts should consider which side in a building block competition will gain an edge as a measure-countermeasure race evolves. For example, AI may enable militaries to field both qualitatively better platforms and quantitatively more platforms, but identifying what the optimal high-low mix looks like requires identifying whether and how the relative advantage between quality and quantity is changing. Focusing only on absolute improvements without assessing relative shifts in cost and capability could cause policymakers to make bad bets on capabilities or concepts that will end up being on the wrong side of the cost-effectiveness curve. To identify rough bounds on the size of potential relative advantages, we use basic analytic models to identify the key variables and trade-offs in the

---

[24] John Stillion and Bryan Clark, *What it Takes to Win: Succeeding in 21st Century Battle Network Competitions*, Center for Strategic and Budgetary Assessments, July 10, 2015.

[25] What distinguishes 21st century "all-domain warfare" from 20th century "combined arms warfare" is how digitized information networks connecting long-range sensors and shooters enable more cross-domain interactions and a wider range of kinetic and non-kinetic effects. The technologies that made combined arms warfare possible in the 20th century included radio communications, aircraft, motorized and mechanized vehicles, improved sensors like radar, and precision-guided munitions. These technologies were often available to both sides, but there were significant differences in how militaries integrated and employed them. See Stephen D. Biddle, *Military Power: Explaining Victory and Defeat in Modern Battle*, Princeton, NJ: Princeton Univ. Press, 2004; and Jonathan B. A. Bailey, "The First World War and the Birth of Modern Warfare," in *The Dynamics of Military Revolution, 1300-2050*, ed. Macgregor Knox and Williamson Murray, New York: Cambridge University Press, 2001, 132–53.

[26] Todd Harrison, "Battle Networks and the Future Force, Part 1: A Framework for Debate," *Center for Strategic and International Studies*, August 5, 2021.

building block competitions. For the quantity-quality competition, we also leverage quantitative models to provide a rough estimate of potential changes to cost-effectiveness.[27]

Using the building block framework, we generate hypotheses about relative advantages in four operational competitions, focusing on long-term military competition between peer adversaries. The analysis also assumes a "mature" AI regime in which the technology is both mature and proliferated.[28] Thus, neither side has an inherent edge simply because they are facing an adversary that has fewer resources or less sophisticated technology to begin with. Because AI technology is not yet mature, these are necessarily tentative hypotheses. The analysis of each building block includes a discussion about the scope conditions for each hypothesis and the conditions under which they might be wrong.

Table 1.2. Hypotheses About the Military Implications of Advanced AI

Building Block Competition	Hypothesized Advantage	Logic
Quality vs. Quantity	Quantity	AI will enable militaries to field large numbers of 'good-enough' autonomous robotic systems at a scale that can overwhelm exquisite systems, while doing so at a cost advantage.
Hider vs. Finder	Contested	AI will enable better finding through improved sensor fusion and fielding more autonomous sensing systems, but it can also enable better deception to offset these gains through "fog-of-war machines" that orchestrate robotic decoys and deceptive tactics that increase the finder's uncertainty. Whether the finder or hider will have an edge will depend heavily on the context.
Centralized vs. Decentralized C2	Mission Command	AI does not change the reasons that mission command as a hybrid C2 has advantages over more centralized or decentralized C2 paradigms, which are rooted in access to the right information rather than in intelligence. But AI could still have important effects for C2 through better coordination and decision-making support across centralized and decentralized elements.
Cyber offense vs. cyber defense	Cyber Defense	AI will facilitate stronger cyber defenses in the long term by addressing some of the limitations with speed, scale, and effectiveness that currently hold the defender back and give the attacker a relative edge.

SOURCE: RAND analysis.

An underlying theme that emerged from the analysis is that some of AI's most significant implications could lie in its effects on the *economics of warfare*. AI has the potential to make mass once again a defining feature of warfare. AI could enable militaries to scale up their operations in ways that resource constraints simply made impractical in the past without having

[27] This method is inspired by the path-breaking work of our late RAND colleague, Lt. Gen (USAF, Ret) Glenn A. Kent. For some of the best examples of this economical mathematical approach, see Glenn A. Kent et al., "Thinking About America's Defense: An Analytical Memoir," RAND Corporation, September 29, 2008.

[28] Andrew Krepinevich, *Maritime Competition in a Mature Precision-Strike Regime*, Washington, D.C.: Center for Strategic and Budgetary Assessments, April 13, 2015.

to recruit and train corresponding numbers of highly-skilled personnel.[29] Coupled with advances in better and lower-cost manufacturing, AI has the potential to shift the cost curve on the scale at which militaries can viably operate. This will matter for direct combat outcomes, such as air-to-air engagements, but it will also shape the other building blocks as well. The hider versus finder building block analysis suggests that the ability to field large networks of ubiquitous sensors and to deploy large arrays of sophisticated physical and electronic decoys will change how deception constrains and enables operations. Additionally, the cyber offense versus cyber defense building block analysis shows how being able to search for vulnerabilities at speed and scale disproportionally benefits the defender because these factors are partly responsible for holding back cyber defense today. In sum, what makes advanced AI valuable is not just that it can replace humans for discrete tasks, but that it can fundamentally change the cost-effectiveness of different asset mixes and operational approaches. Together, changes to the underlying cost-effectiveness of building and employing forces could eventually constitute an RMA that will significantly disadvantage militaries that fail to adapt.

The remainder of this working paper proceeds as follows. The next four chapters present the analysis of the four operational building blocks. Chapter 6 explores the implications of a first-mover advantage where one side has AI-enabled military capabilities before its adversary develops them. Chapter 7 then steps back to assess the collective implications for future warfare and highlights remaining uncertainties and directions for further research.

[29] See, also, Jacquelyn Schneider and Julia Macdonald, "Looking back to look forward: Autonomous systems, military revolutions, and the importance of cost," *Journal of Strategic Studies*, Vol. 47, No. 2, 2024, pp. 162–84.

Chapter 2. Quality vs. Quantity

How will advances in AI influence the relative advantage between *quantity* and *quality* in military operations?[30] A recurring theme in military history is that new technologies have sometimes allowed one side to offset a numerical disadvantage by creating a qualitative edge, such as the British longbowmen at Agincourt and U.S. breakthroughs in stealth and precision strike networks during the Second Offset.[31] The U.S. military's force planning in recent decades has relied heavily on building this type of qualitative edge by fielding fewer but more capable platforms. The extent of this emphasis on quality over quantity has raised concerns that the U.S. military may struggle against an adversary that prioritizes mass in its own force structure because of the adage that "quantity has a quality of its own."[32]

We argue that AI could produce greater relative advantages for quantity over quality. This stands in contrast to most technological innovations in the past, which tended to benefit quality. The underlying intuition for this hypothesis is that improved autonomy will make it easier to field N times more platforms than to make an individual platform N times more effective. While AI will enable improvements in both quantity and quality, it is unlikely to offer large enough qualitative advantages for a small number of exquisite forces to counterbalance a mass of good-enough forces. Critically, AI may enable militaries to leverage quantity in ways that have been unaffordable in the past.

Given that AI will enable advances in both quantity and quality, the key analytical challenge is to develop a framework to adjudicate which of the two will likely benefit more. The goal is to test the hypothesis that AI will give quantity an edge over quality in a principled way that produces clear predictions, identifies conditions under which the hypothesis is wrong, and enables other researchers to replicate the analysis. We use Lanchester equations as the framework for testing the hypothesis because the equations provide a clearly structured and well-established method for exploring trade-offs between quantity and quality.[33] Additionally, our analysis focuses on the quality-quantity competition in the context of air-to-air combat, though

[30] For other recent work on the competition between quality and quantity, see: Michael Horowitz, "Battle of Precise Mass," *Foreign Affairs*, Vol. 103, No. 6, 2024; and Marcel Plichta and Ash Rossiter, "A one-way attack drone revolution? Affordable mass precision in modern conflict," *Journal of Strategic Studies*, Vol. 47, No. 6, November 9, 2024, pp. 1001–31.

[31] Michael Howard, *War in European History*, Oxford, UK: Oxford University Press, 1976; Krepinevich, 2023.

[32] See, for example, Mackenzie Eaglen, "The Bias for Capability over Capacity has Created a Brittle Force," *War on the Rocks*, November 17, 2022.

[33] For a useful introduction, see: John Lepingwell, "The Laws of Combat? Lanchester Reexamined," *International Security*, Vol. 12, No. 1, Summer 1987, pp. 89–134; and Alan Washburn, *Lanchester Systems*, Monterey, CA: Naval Postgraduate School, March 1, 2000.

we anticipate the trend toward the growing importance of mass extends to other operations and domains as well, for reasons that we will explain.

How AI Can Help Both Quantity and Quality in Air Combat

We hypothesize that advanced AI could enable militaries to significantly increase the number of combat aircraft in their force structures by fielding AI-enabled autonomous aircraft. This development could enable smaller and lighter aircraft that are less expensive.[34] The extent to which these systems are cheaper depends partly on their requirements, such as the quality of their sensing systems, their range, and their payload capacity. If militaries wanted to develop extremely long range, low observable, highly maneuverable unmanned aircraft with large payloads and cutting-edge radar systems, then the cost savings from removing a human pilot from the cockpit may be small. This means that militaries would need to think differently about capability requirements in force design to realize potential cost savings. Put differently, states would need to choose to use AI to build numerous cheaper autonomous fighters rather than exquisite expensive ones. Policymakers may be more willing to reduce demanding requirements that have historically driven up costs partly because losing aircraft no longer means the loss of human life. On average, the unit price will likely decrease if militaries embrace this new approach, and across economies of scale this would produce a cost advantage for building larger numbers of aircraft.

There are several other mechanisms through which AI could reduce resource requirements and enable militaries to increase the size of their force structure. AI would also reduce human capital requirements for skilled pilots, which is a major determinant of effectiveness and a limiting factor for many militaries.[35] Without having to continuously fly aircraft to train human pilots and keep their skills sharp, aircraft maintenance requirements would likely decrease significantly. These platforms could start to look more like missiles, where some are flown each year to test the quality of the inventory but do not require regular use outside of combat. Realizing this scale in practice would require improvements to the U.S. defense industrial base, so AI's potential to generate large efficiency gains in production as outlined in Chapter 1 would be an important enabler of generating the kind of mass this chapter envisions.

At the same time, AI will likely improve the quality of aircraft in tactical engagements. In aerial combat, the side that can fuse, analyze, and act effectively on data faster than its opponent gains a significant advantage. If AI can make decisions better and faster than human pilots, it could gain a decisional advantage over manned aircraft that translates into a combat advantage. Removing the physical constraints of human pilots could also enable aircraft to operate at longer

[34] See, for example, Gregory C. Allen and Isaac Goldston, "Updating Augustine's Law: Fighter Aircraft Cost Growth in the Age of AI and Autonomy," Center for Strategies and International Studies, December 19, 2024.

[35] Rachel Cohen, "Perennial pilot paucity puts Air Force in precarious position," *Air Force Times*, March 3, 2023.

ranges without exhausting pilots and reducing their performance, as well as enable tactics that would be too dangerous or strenuous for human pilots but do not strain AI systems.[36] For example, autonomous aircraft might become more survivable if they can react faster and make more aggressive maneuvers than humans can in order to evade a missile.

How Much Quality is Enough to Overcome Greater Quantity?

Lanchester equations describe battles whose outcome depends on attrition. Specifically, they are differential equations for modeling battles of attrition between two sides that last until one side loses all its forces. There are several variants of these equations, and we use the "Lanchester Square Law" because it best matches the character of modern warfare, especially in its core assumption that multiple units can concentrate their fire on one target.[37] Other key assumptions of the Lanchester Square Law equations include:

- Homogenous composition of forces on each side (e.g., fighters face fighters, not bombers).
- A battle's outcome is based on attrition (that is, the winner is the side that grinds down and eliminates the loser's entire force; battle outcomes are not based on maneuver or operational art).
- Both sides can see and engage all the other's forces, so the battlefield is a target rich environment where it is possible to concentrate fire on any given target.[38]
- Both sides engage each other simultaneously, so there is no "first mover" advantage. This assumption implies that neither side has such a large advantage in beyond visual range combat that their fighters can simply engage the enemy before its fighters have a chance to fight back.[39] In other words, neither side gets to carry out a first strike.[40]
- Units in the battle retain the "will to fight" and will keep fighting despite poor odds of survival.

---

[36] Scharre, 2023.

[37] There are two main forms of the equations – one where each unit can attack multiple targets and receive fire from multiple directions (which Lanchester called "collective combat") and another where each unit can only attack one target and receive damage from one target (which Lanchester called "individual combat"). We use the "collective combat" model, commonly known as the "Lanchester Square Law." Lepingwell, 1987.

[38] For modeling simplicity, we assume that both sides are equivalently skilled in the use of jamming and decoys such that the effect nets out and both sides are equally able to find and engage real targets. Conceivably, AI could help one side gain an advantage in deception that changes this dynamic.

[39] If AI could dramatically increase the effectiveness of beyond visual range combat, that would be one way to increase lethality that falls outside the model's scope. However, key limitations on the range of engagements are more likely to be the physical limits of sensors (e.g., line of sight to the horizon) and rules of engagement (e.g., ensuring that a fighter is about to engage an enemy aircraft, not a friendly or civilian aircraft).

[40] Lanchester equations have historically been better suited to air combat because of its melee character; fighters have to place themselves at risk to attack the enemy. By way of contrast, naval combat in the aviation and missile ages has depended much more on striking first effectively. See: Wayne P. Hughes, *Fleet Tactics: Theory and Practice,* Annapolis, Md: Naval Inst. Press, 1986.

AI might enable autonomous units to engage in combat in ways that more closely match the assumptions of Lanchester equations than human combat has (Table 2.1). Analysts have often used Lanchester equations, but their simplifying assumptions mean they are often poorly suited to reflecting important combat dynamics. Attempts to validate Lanchester models with historical battle data have been mixed at best.[41] However, battle between autonomous robotic agents might more closely match the key assumptions behind the Lanchester equations, making them a useful tool for considering the implications of AI-enabled robotic warfare. To be clear, our argument is not that every assumption behind the Lanchester Square Law will always hold but that the Law provides a useful starting point for assessing tradeoffs between quantity and quality in robotic warfare.

Table 2.1. Lanchester Assumptions: Human Battle vs. Robotic Battle

Key assumptions of Lanchester Equations	Does assumption apply to human battle?	Would this assumption apply to autonomous robotic air battle?
All forces can engage each other	Sometimes	Sometimes
Homogenous agents	Sometimes	Sometimes
Agents willing to fight despite poor odds of survival	Rare	Yes
Battle outcome determined by raw attrition	Rare	Sometimes

Note: Homogenous agents in the Lanchester Equations refers to symmetrical capabilities, like fighters facing fighters. AI would likely produce more homogenous agents at the individual-level by reducing human-level variation in performance and decision-making.

In its basic form, the Lanchester Square Model describes one force (Force A) with quantity x and an effectiveness of α. Effectiveness is measured in terms of the quantity of enemy units killed by each x per unit time t. Force A is in combat with a second force (Force B) with quantity y and effectiveness β. The derivatives of x and y with respect to t are:[42]

$$\dot{x} = -\beta y$$
$$\dot{y} = -\alpha x.$$

[41] For example, Engel, 1954 found a reasonable fit for the Battle of Iwo Jima if one included American reinforcements but other efforts have failed to find good fits for Lanchester equations to the Battle of the Ardennes in 1940 and the Battle of Kursk in 1943. J.H. Engel, "A Verification of Lanchester's Law," *Journal of the Operations Research Society of America*, Vol 2, No. 2, May 1954, pp. 163–71.; Jerome Bracken, "Lanchester Models Of The Ardennes Campaign," *Naval Research Logistics*, Vol. 42, Issue 4, 1995, pp. 559–57; Turker Turkes, *Fitting Lanchester and Other Equations to the Battle of Kursk Data*, thesis, Naval Postgraduate School, 2000.

[42] Washburn, 2000.

To summarize, in the simplified world of Lanchester combat, there are *only two ways to improve the combat power* of one's force: *increase quantity* or *increase lethality*.

- *Increase lethality*. This would require breakthroughs in engineering to increase the number of potential kills that finite platforms can produce, which makes it less plausible that AI by itself will enable a revolution in lethality. Additionally, if qualitative improvements increase one side's survivability, that will decrease the other side's relative lethality. However, the nature of Lanchester combat means that chasing lethality offers less leverage to a force planner than pursuing quantity.
- *Increase quantity*. AI might enable levels of autonomy that is good enough that, when combined with procurement in sufficient numbers, can prevail even over opponents that are more lethal on a per-unit basis. For example, the U.S. Air Force (USAF) would not need to field a collaborative combat aircraft (CCA) with AI that is as lethal as a manned Chinese fighter in order to realize an advantage. If the USAF could field twice as many CCA in a battle than Chinese fighters, then the CCAs could win so long as their lethality is not less than 1/4 of the Chinese fighters.

If the Lanchester Square Law accurately captures key dynamics for air-to-air engagements enabled by new AI breakthroughs, then quality will face an uphill battle against quantity. One of the benefits of these equations is that they clearly identify the frontier at which advances in quality could in fact overcome disadvantages in quantity. We can visualize how much of a quality advantage would be necessary to overcome a given numerical disadvantage (obviously, if one side has both a quality and quantity advantage then that side wins). Figure 2.1 graphs the Lanchester Square Law. In the blue region, the more effective side in terms of lethality wins (for example, a side with a 9:1 advantage in lethality can win if it is outnumbered 2:1 as shown by point A). In the red region, the more numerous side wins (for example, a side with a 4:1 quantitative advantage can win even if its adversary has a 2:1 advantage in lethality as shown by point B).

Figure 2.1. Quality vs. Quantity Break-Even Points (Using the Lanchester Square Law)

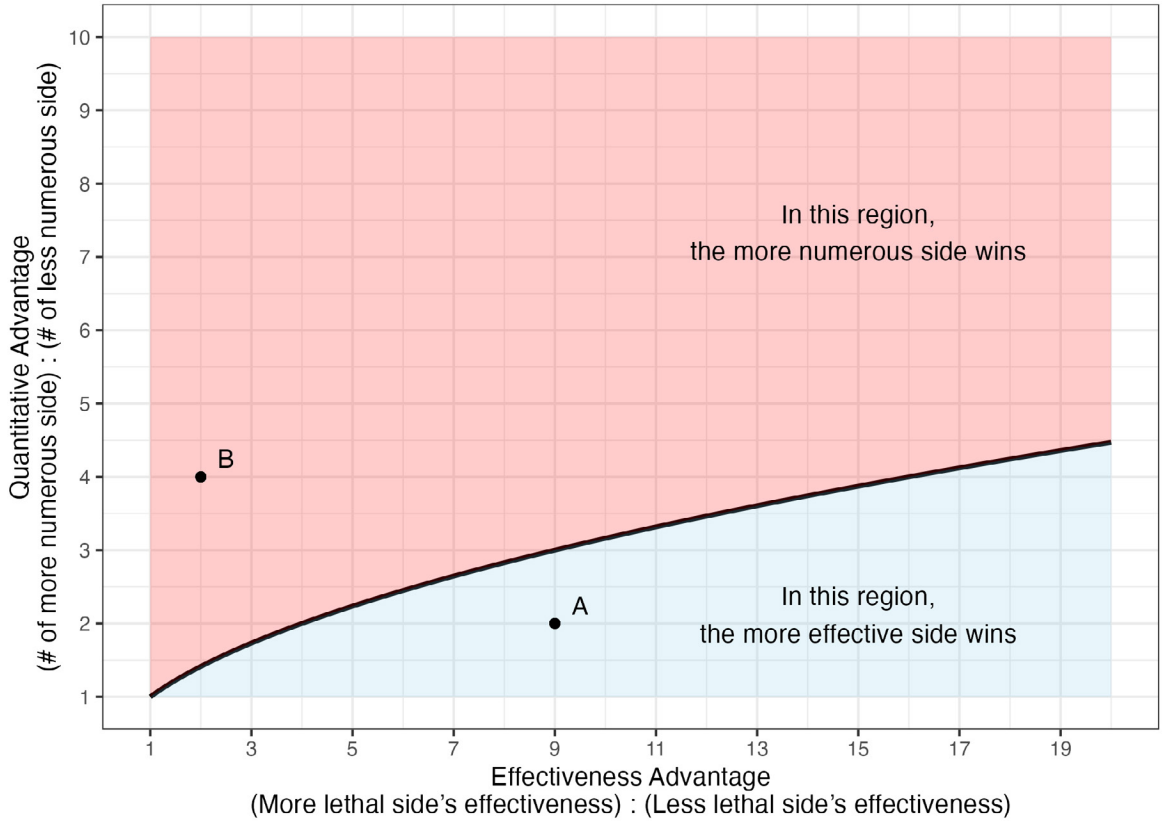

Historical examples can provide a sense of what a realistic upper bound on lethality advantages might be absent a radical breakthrough in physics that could enable completely new modes of combat. Building on Figure 2.1's model, Figure 2.2 adds a historical example from air-to-air combat during World War II. Late in the war, Germany deployed the Me 262, the world's first jet fighter (U.S. aircraft, such as the P-51D fighter, still used piston engines at the time). The Luftwaffe generally staffed its Me 262 units with extremely experienced German aces. A battle between the Me 262 and the P-51D is a useful historical illustration of a breakthrough technology giving one side a significant lethality advantage. Because of these advantages in superior technology and pilot skill, historical data on air-to-air engagements between the Me 262 and the P-51D suggests that the Me 262 likely had around a 9:1 lethality advantage over the P-51D.[43] Following the Lanchester Square Law, that suggests that the Me 262 fighter could

[43] To estimate the Me 262's lethality advantage, we draw on two days of intense aerial combat in 1945 that represent the peak of the Me 262's use in the war. John Foreman and S. E. Harvey, *The Messerschmitt Combat Diary Me.262*, Crecy Pub, 2nd ed., 1995, pp. 176-181, 236-244, 339, 344-345. On March 18, the United States generated 600 fighter sorties compared to 37 German sorties of Me 262s. Germany achieved 21 kills against U.S. fighters, and the United States killed 6 Me 262s. On April 10, the United States generated 815 fighter sorties compared to 55 German sorties of Me 262s. Germany achieved 4 kills against U.S. fighters, while the Germans lost 33 Me 262s. Combining these days, the United States had 1,415 fighter sorties, Germany had 92 fighter sorties, the United States had 39

overcome being outnumbered by 3:1. The Me 262 example is a reasonable upper bound on the maximum advantage that superior equipment might provide against a skilled adversary. If advantages in relative lethality are significantly greater than 9:1, then the enemy is unlikely to be a peer competitor, as was the case in Israeli's lopsided air-to-air victories against much weaker Arab militaries in the past.[44] If 9:1 provides a reasonable upper bound on the advantages next-generation technologies provide against a peer adversary, then the best that lethality improvements might offer is the ability to fight outnumbered 3:1. Therefore, if AI could enable more than a 3x increase in the quantity of systems that are at least $1/9^{th}$ as lethal as what the adversary fields, then the side with the larger force would prevail.

fighter kills, and Germany had 25 fighter kills. Determining the lethality advantage for the Lanchester Square Law requires accounting for the number of fighters involved in engagements. Looking solely at loss-exchange ratios conflates advantages derived from quantitative advantages with advantages from qualitative advantages in lethality. Adjusting for the sortie ratio, these statistics suggest the Me 262 had a lethality advantage of 9.86 (1,415/92 * 25/39 = 9.86). In our analysis, we round down from 9.86 to 9 because of the illustrative simplicity of using a perfect square and because the difference between the two numbers is too small to change our conclusions. These statistics are intended to provide an illustrative upper bound, but they are necessarily imperfect and different approaches to counting could lead to slightly different estimates. These were not truly "pure" fighter vs. fighter duels between the Me 262 and the P-51. The Me 262 was also targeting U.S. bombers, but the imperative to hunt bombers is what forced German fighters to engage U.S. fighters in the first place. Additionally, on March 18, only 10 of the 21 U.S. fighter losses were P-51s. Excluding the other fighter types, the Me 262's relative lethality against the P-51 would fall to 5.52. By contrast, on April 10, U.S. fighters destroyed three Me 262s in a strafing run while they were still on the ground. Excluding those losses would increase the Me 262's relative lethality to 10.68. Making both adjustments would lead to an Me 262 lethality advantage of around 6:1, which would fall even more if the data were available to reduce the U.S. sortie rate on March 18 to only include P-51s and no other fighter types. The estimate of a 9:1 advantage is therefore well-grounded empirically as a strong upper bound on what the Me 262 might have had over the P-51.The overall results of the analysis are not sensitive to modest adjustments to this upper bound. For excellent historical analysis on trends in fighter qualitative advantages that supports the assessment that a nine-to-one advantage was still an upper bound in subsequent decades, see: John Stillion, *Trends in Air-to-Air Combat: Implications for Future Air Superiority*, Center for Strategic and Budgetary Assessments, April 14, 2015.

[44] Kenneth Pollack, *Arabs at War: Military Effectiveness, 1948-1991*, University of Nebraska Press, 2002.

Figure 2.2. Quality vs. Quantity Break-Even Points with Historical Example

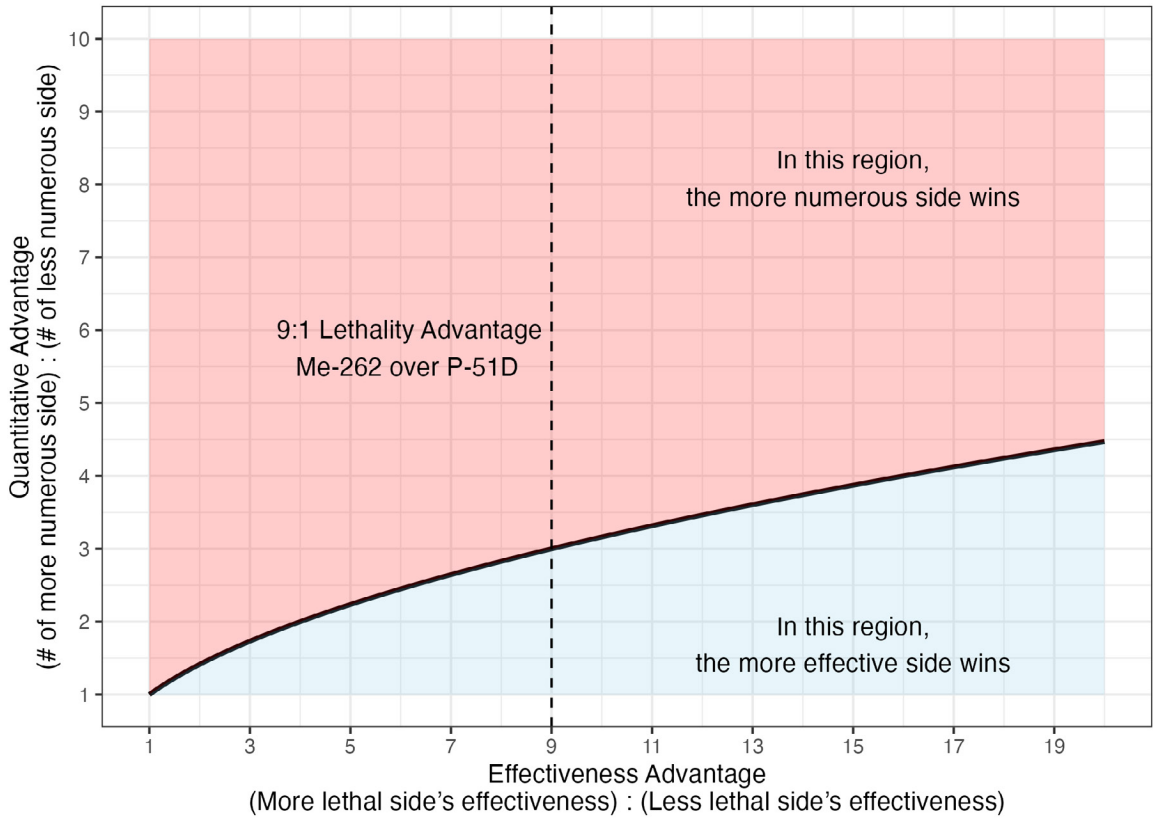

This supports the hypothesis that AI could give quantity an edge over quality. AI-enabled autonomy that facilitates fielding large numbers of systems at a reasonable quality is more valuable than chasing extreme lethality improvements. This implies that an investment strategy to create a mass of good-enough AI robots could beat an opponent with "superhuman" performance; a force of 401 drones would defeat a force of 200 drones, even if the 200 drones were 4x as lethal. Or, put differently, one should choose a force of 950 autonomous P-51s over 300 autonomous Me 262s, so long as the autonomous P-51s were at least 1/9 as lethal as the autonomous Me 262s.

What are Potential Limiting Factors on Embracing Mass?

Caveat 1: Cost

If the Lanchester Square Law suggests that quantity has such a strong advantage over quality, why have more militaries in the past not focused their force structures more squarely on mass? One potential reason is that the assumptions of the Lanchester Square Law struggle to map onto the dynamics of human combat, causing them to overstate the advantages of quantity. But if the

assumptions more closely resemble robotic warfare, the relative advantages of quantity may increase in the future, and militaries may reorient their force structures as a result.

Another potential reason is that generating enough mass is simply expensive. If the poorer country is generally the one trying to rely on sheer numbers to offset the qualitative advantages of an opponent, then it must generate enough force to counterbalance its lethality disadvantage. Relying on sheer numbers can actually be quite expensive, especially for systems like fighter aircraft. This is especially true because the break-even analysis above highlights that a mass strategy still requires fielding units that have a sufficient baseline of lethality, which drives up cost. The logistical burdens of supporting an exponentially larger force structure will also increase costs, though further analysis is necessary to compare these costs to the requirements that a smaller number of exquisite aircraft create.

This cost consideration suggests that quantity may have a combat advantage at the tactical level but that generating a much larger force structure may be cost ineffective and financially unsustainable at the strategic level. Even if the unit price is cheaper, the cost of the total force structure may be more expensive. In a long-term competition where neither side can afford to dramatically outspend the other, investing in mass may be impractical.

Our historical example can illustrate these cost dynamics. Fielding 950 P-51Ds may be more expensive than fielding 300 Me 262s. Based on the capabilities of those platforms, we estimated that the cost ratio for an Me 262 to a P-51D is around 1.5.[45] That is, building an aircraft as capable as the Me 262 would cost 1.5 times more than building an aircraft as capable as a P-51D based on their characteristics like weight and speed. If the U.S. P-51Ds need to outnumber the German Me 262s by more than 3:1 to offset the 9:1 German lethality advantage (Figure 2.2), then a P-51D would need to be less than 1/3 the price of an Me 262 to come out cost neutral. But the estimated price is closer to 1/1.5. As a result, fielding the required number of P-51Ds to prevail over the Me 262s would be more expensive.[46] For the United States in World War II,

---

[45] This uses cost-estimating relationship equations. Cost-estimating relationships use attributes of a system to estimate how much it will cost. Cost-estimating relationships are developed through statistical analysis of past production. The ratio here is not an exact historical estimate of the unit price for each. It is a comparison of the costs required to build aircraft as capable as the Me 262 and the P-51D based on their performance characteristics, holding all else equal. We used a RAND cost-estimating relationship from 1973 in which total cost is the sum of the engineering costs, development costs, tooling costs, flight test operations costs, manufacturing labor costs, materials costs, and quality control costs. G.S. Levenson et al., *"Cost-Estimating Relationships for Aircraft Airframes*", RAND Corporation, R-761/1-ABR, January 1, 1972, pp. 53-56. We held constant factors that were unrelated to aircraft performance, such as labor costs, the total number of airframes produced, the number of test airframes, and the monthly production rate. The two key aircraft parameters we varied were the maximum speed and weight of the aircraft. We clarify above that this is not a historical cost estimate of the actual cost difference between the Me 262 and the P-51D because the United States and Germany would have had differences in production costs, such as the cost of labor, that were unrelated to aircraft performance. But the cost-estimating relationship is nevertheless useful in providing a notional comparison of aircraft that would have had comparable differences in capabilities similar to that between the Me 262 and the P-51D.

[46] For example, to win a battle with 10 Me 262s, the United States would have to commit more than 30 P-51Ds and the U.S. forces committed would cost more than twice as much as Germany committed [(30/10) * (1/1.48) = 3/1.48 ~ 2].

these kinds of cost inefficiencies were acceptable because the United States could out-produce Germany. In a competition between two countries with equivalently strong defense industrial base capacity, that might not be the case.

Figure 2.3 overlays the World War II example to the break-even analysis from the previous figures. The horizontal line shows the cost ratio for an Me 262 relative to a P-51D. In the red region, the quantitatively larger force of P-51Ds still prevails over the smaller force of Me 262s, but including considerations of cost reveals that it does so at a cost disadvantage. That is, the United States could win the battle, but it would pay disproportionately more for the victory. Only in the relatively small purple region would the United States both win and do so at a cost advantage over Germany.

Figure 2.3: Cost Calculations and the Affordability of Quantitative Advantages

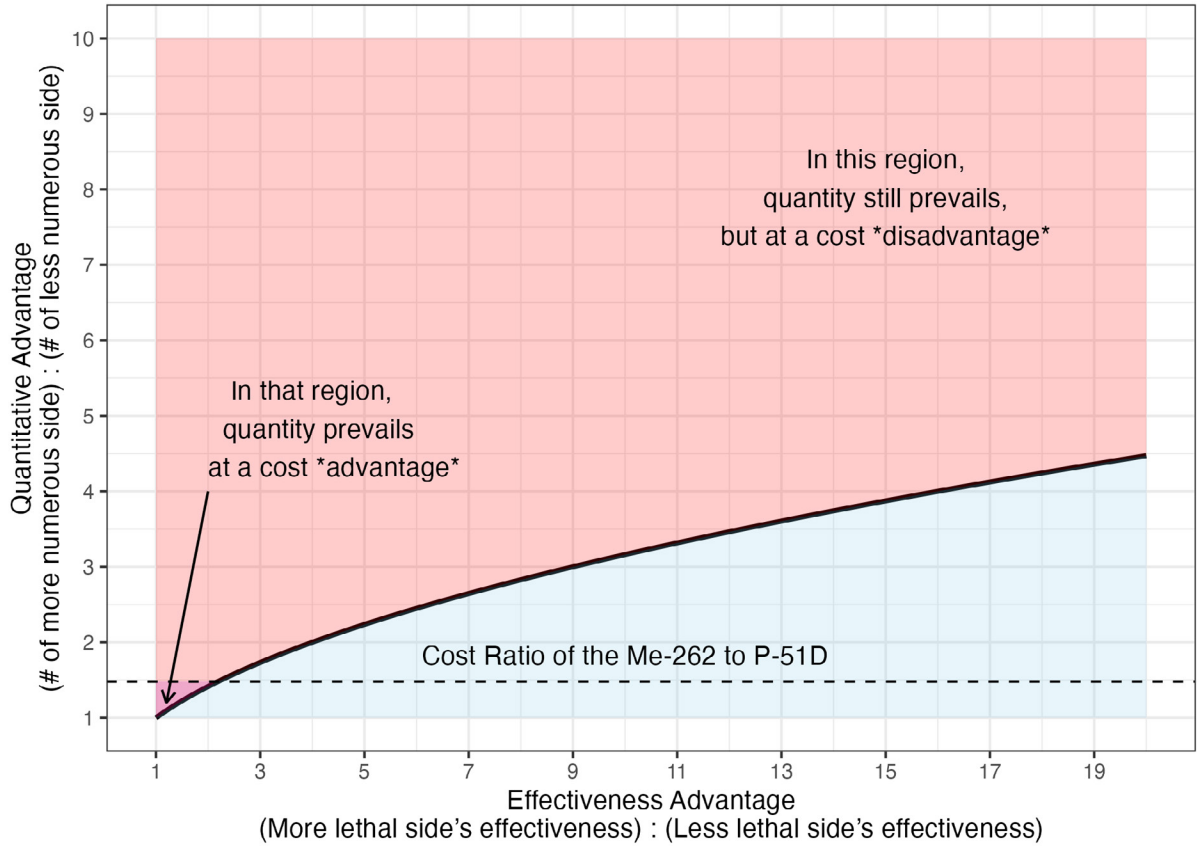

However, AI may change the dynamics of this cost competition. It may allow one side to field dramatically cheaper platforms without a correspondingly dramatic reduction in combat capability. Figure 2.4 visualizes the same break-even analysis as above, but it uses the cost ratio from a forward-looking rather than a historical case. It plots the estimated cost ratio of a Chinese J-20 manned fighter compared to a notional autonomous robotic fighter that has characteristics

comparable to the XQ-58 Valkyrie, which is an experimental U.S. unmanned combat aerial vehicle (UCAV). One J-20 is an estimated 4.71 times as expensive as a notional XQ-58-like robotic fighter.[47] Figure 2.4 illustrates that the hypothesized cost reduction that advances in AI and robotics enable significantly expands the purple region where quantity can prevail at a cost advantage. The most dramatic impact of AI might therefore be enabling cost-effective mass in a way previously not possible.

Figure 2.4: How Autonomy may Change the Cost of Relying on Quantity

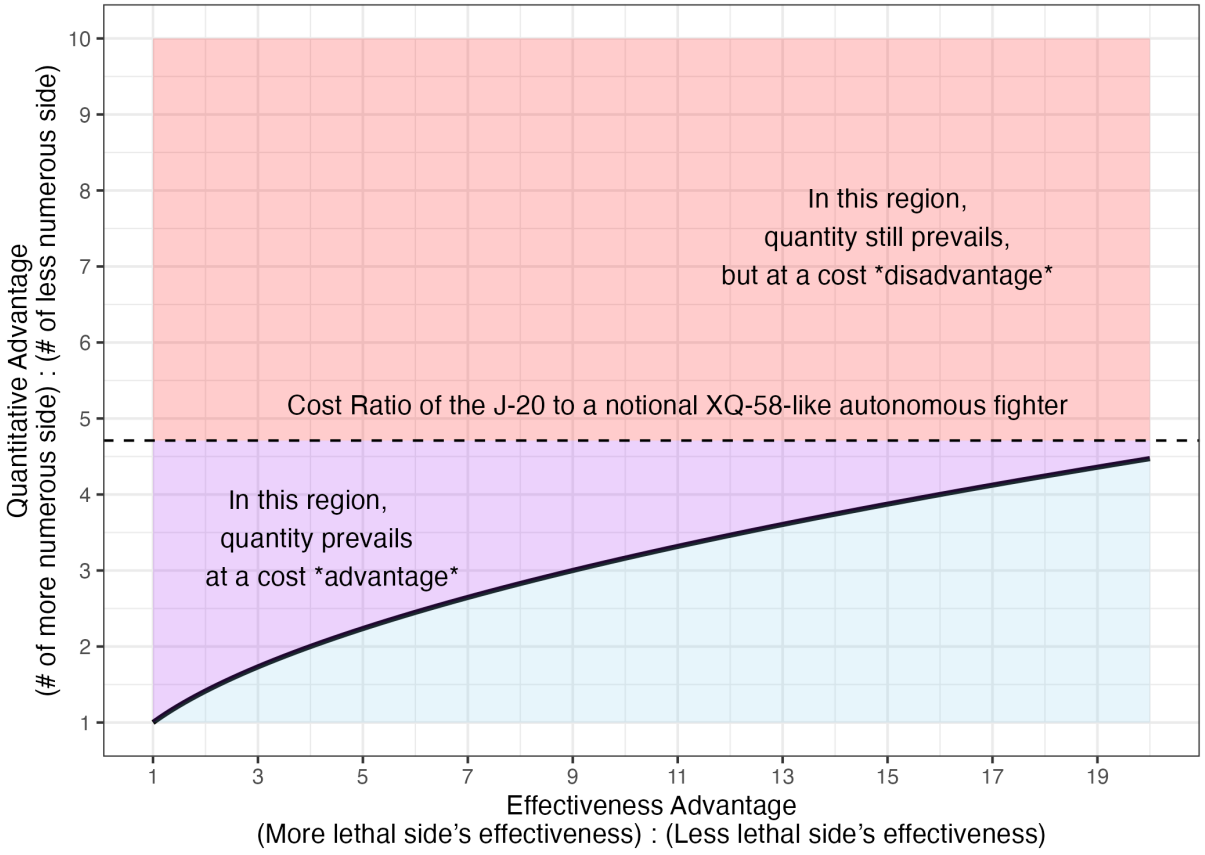

[47] The parameters and their coefficients are a dummy variable for a remotely piloted vehicle (0.299), the log of empty weight (0.651), a dummy variable for stealth (0.747), the log of the production rate (–0.345), the year of the first delivery (0.006), an indicator for the mission type (0.339 for fighters), and an intercept (–13.123). The dependent variable is the log of the average unit fly-away costs of the first 100 aircraft. For our purposes, we were interested in comparing the relative costs of two different types of aircraft, so the exact form the costs take in absolute terms is less important, as are such variables as the year of the first delivery as long as the values for them are held constant across the two aircraft of interest to focus the comparison on differences in capability attributes. Thomas Light, Fred Timson, and Obaid Younossi, *Parametric Flyaway Cost Estimating Relationships for Manned and Remotely Piloted Military Aircraft*, RAND Corporation, RR-1618/1-AF, 2024, p. 11. The estimate of the empty weight for the J-20 (42,747 pounds) came from Janes, and the estimate of the empty weight for the XQ-58 (2,500 pounds) came from a factsheet produced by the manufacturer Kratos. "CAC J-20," February 3, 2021, *Janes*; "Valkyrie XQ-58A," Kratos, No Date.

Caveat 2: Geography

Geography is another variable that could qualify our findings. Distance matters for each side's ability to bring mass to bear in each battle. The farther one side's forces must fly from an air base to a battle, the fewer sorties (missions) a fixed force will be able to generate.[48] It can also require fielding more capable and more expensive aircraft that have greater range. The closer one side's air bases are to a battle, the more sorties it can generate and the more viable it is to build cheaper, shorter-range fighters. Asymmetries in geographic context can make the more distant side build larger, more expensive aircraft and force it to build more of them to reach quantitative parity with the geographically closer side over the battlefield.

Drawing on the U.S.-China context, Figure 2.5 provides a basic illustration of how geography imposes its own constraints on mass. The closer that air-to-air engagements occur to China, the farther they occur from U.S. air bases in the region, such as those in Japan or Guam. Right off China's shores, China has a major advantage in generating sorties because its aircraft do not need to fly for very long, so it can keep larger numbers of aircraft on station for longer periods of time. All else equal, each side operating the same number of aircraft from distance distances would give China a 2.7:1 force advantage over U.S. fighters along its coast, with the size of this advantage diminishing as air-to-air engagements take place farther from China. If the United States wanted to achieve quantitative parity in air-to-air engagements close to China, it would already need to use a larger force structure than China (or else leverage surprise). If both sides developed robotic air superiority aircraft, all else equal the U.S. aircraft would need to be more expensive to give them greater range if they needed to operate from distance bases like Guam (or the United States would need to use transport aircraft to get shorter-range aircraft closer to the fight).

[48] John Stillion and David Orletsky, *Airbase Vulnerability to Conventional Cruise-Missile and Ballistic Missile Attacks*, Santa Monica, CA: RAND Corporation, 1999, pp. 81-84.

Figure 2.5: Asymmetries in Distance Increase the Quantity of Aircraft Required

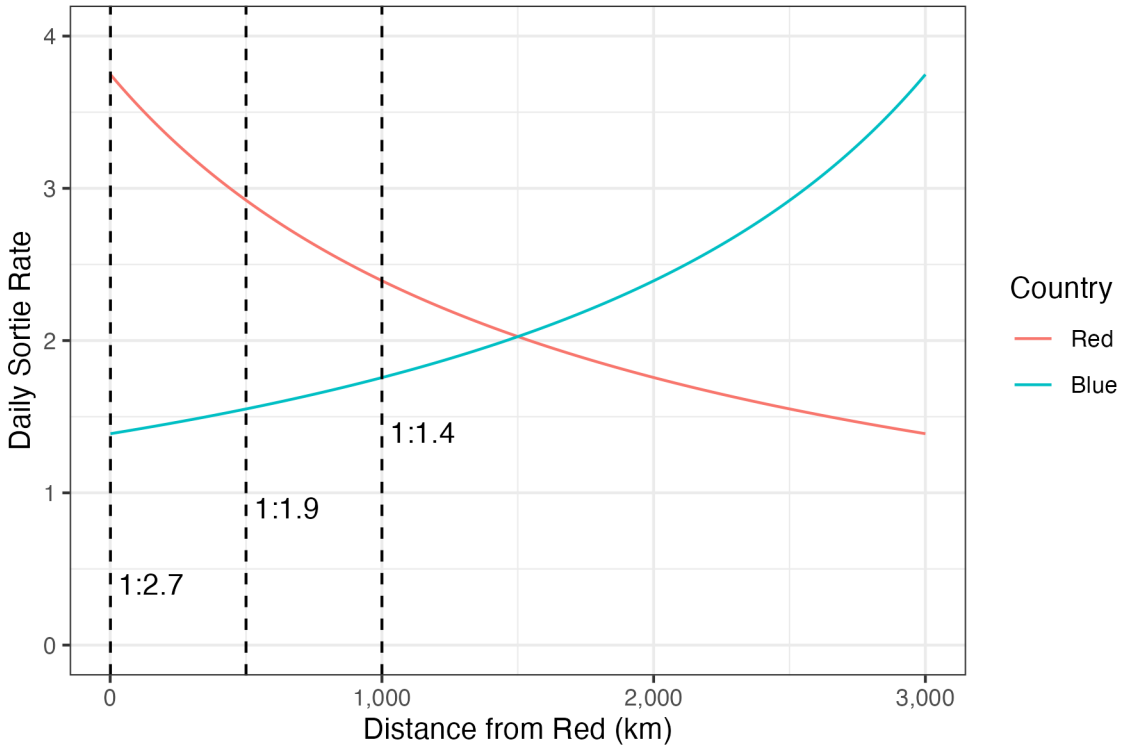

Note: Ratios in Figure 2.5 reflect the sortie advantage of Red at a given distance from China.

In the context of a war over Taiwan, if the United States follows its legacy approach to power projection, it could be operating from a small number of air bases that would be farther from the battlefield than China's air bases. If the United States maintained this approach, then the geographic asymmetry would prevent the United States from achieving the full potential of the cost advantages for autonomous fighters that we discussed earlier. Smaller autonomous fighters could still offer advantages, but not as much as postulated for the XQ-58 example. If the United States does not change its approach to power projection, then there is reason to moderate expectations that autonomy will create a large U.S. cost advantage over China.

If, however, the United States explored a new approach to power projection that relied on large quantities of runway-independent UCAVs, then the United States might be able to realize the aforementioned advantages discussed by basing these UCAVs in multiple distributed sites across Taiwan, southwest Japan, and the northern Philippines.[49] This concept would avoid the missile threat to U.S. and allied air bases in the region by detaching at least some airpower from

[49] See, for example, Thomas Hamilton and David A. Ochmanek, *Operating Low-Cost, Reusable Unmanned Aerial Vehicles in Contested Environments: Preliminary Evaluation of Operational Concepts*, RAND Corporation, RR-4407-AF, May 11, 2020.

fixed operating locations that Chinese cruise and ballistic missiles could easily target and would then rely on massed autonomous UCAVs to accomplish combat missions.[50] The quantity of these UCAVs could overcome the poorer effectiveness of each UCAV compared with a manned Chinese fighter.

Caveat 3: Generalizability to other domains

While the analysis here focuses on air-to-air engagements, quantity may also have growing advantages in other domains as well. As China strengths its space and long-range strike capabilities, large U.S. surface combatants could become increasingly vulnerable at greater and greater ranges.[51] This is consistent with emerging naval concepts of distributed maritime operations and "distributed fleet combat," where high-value ships are distributed across the battlespace to reduce their exposure concentrated missile salvos.[52] These trends will likely reinforce long-standing calls for the United States to increase its emphasis on smaller and unmanned surface ships and undersea capabilities relative to the current emphasis on a limited number of large surface combatants, especially carriers.[53]

Identifying the likely cost curve and force structure mix in the naval domain is an important area for future research. One of the 20th century's preeminent naval theoreticians, Captain Wayne Hughes, noted that the Lanchester equation's assumptions like the importance of concentrating firepower do have some relevance to naval warfare, but a key difference is that naval warfare in the missile age places a much greater emphasis on the timing of salvo pulses.[54] That is, there are large advantages of firing first in naval warfare that the Lanchester Equations do not fully capture. Additionally, there are other aspects of naval warfare that require different kinds of analysis, such as how having fewer or no crew on board changes the platform's ability to conduct maintenance and battle damage control. These caveats are not definitive reasons that mass will not also have a growing advantage over exquisite systems in naval warfare, but they are reasons to be cautious about extrapolating conclusions from the air-to-air engagements example shown here to naval operations without additional analysis.

---

[50] The dispersed launch and recovery elements operating the runway-independent UCAVs could employ the deception methods discussed in Chapter 3 to improve their survivability.

[51] Audrey Decker, "Chinese satellites are breaking the US 'monopoly' on long-range targeting," *Defense One*, May 2, 2024.

[52] Jeffrey R. Cares, Anthony Cowden, and Scott H. Swift, *Fighting the Fleet: Operational Art and Modern Fleet Combat*, Blue & Gold Professional Library, Annapolis, Maryland: Naval Institute Press, 2021, pp. 86-87.

[53] For an analysis of the implications of unmanned systems for the future fleet, see: Scott Savitz and Amanda Perez, *Could the U.S. Navy Fleet of the Mid-21st Century Include Large Uncrewed Vehicles?,* RAND Corporation, PE-A2854-1, January 2025.

[54] Wayne P. Hughes, *Fleet Tactics: Theory and Practice*, Annapolis, Md: Naval Inst. Press, 1986, pp. 37-38 and 271-272.

Conclusion

We hypothesize that AI could give quantity a relative advantage over quality. The key reason is that the possibility of fielding large numbers of "good-enough" small robotic fighters at scale would improve the cost-effectiveness of relying on mass over quality. Our analysis focuses on the context of air warfare, but this trend could apply to other domains as well.

The growing advantages of quantity over quality represent a reversion in some ways to long-term historical trends. For most of history, militaries have viewed mass as a key factor in giving one side a battlefield edge over the other.[55] Militaries could sometimes offset quantitative disadvantages if they could leverage technological innovations faster or better than their opponents. Famous examples include the British longbowmen at Agincourt and British colonial troops armed with the Maxim gun in colonial conflicts.[56] These types of qualitative advantages had generally been short-lived and did not change the larger preference for quantity over quality once a technology had diffused. Toward the end of the Cold War, however, the United States pioneered advances in stealth and precision-strike networks that created overwhelming and long-lasting qualitative advantages. While exquisite capabilities like stealth aircraft significantly drive up cost, their performance advantages have made it more cost-effective to rely on fewer but better platforms and weapons for the past 50 years. The analysis here suggests that improvements in robotics and autonomy are starting to erode these advantages.[57] As lower-end systems become more capable without a correspondingly large rise in cost, relying on a small number of qualitatively superior platforms will become less cost-effective in turn. While exquisite platforms certainly have a role in future battlefields, the optimal mix between quality and quantity could increasingly favor the latter.

[55] Horowitz, 2024.

[56] Scharre, 2018, p. 38.

[57] See, also, Horowitz, 2024.

Chapter 3. Hider vs. Finder

There is a longstanding sense among analysts and policymakers that technological advances will eventually create a "transparent battlefield."[58] Some believe that the mass employment of unmanned sensing systems coupled with AI information fusion systems might finally realize this vision. The head of the U.S. Army Futures Command predicted that by 2040 the "ability to hide, which is fundamental to how we fight, is now impossible."[59] We argue that sweeping predictions of a transparent battlefield overlook the hider's ability to benefit from AI.[60] If militaries invest in robotic decoys and AI-enabled "fog of war" machines, which are AI tools that orchestrate elaborate deception campaigns designed to maximize the finder's confusion and uncertainty, then finders will once again discover that new technologies have failed to lift the fog that is endemic to warfare.[61]

Advances in AI will enhance the capabilities of both "hiders" and "finders" in absolute terms. Militaries can use AI not only to enhance "finder" activities, like automated target recognition, but also "hider" activities, such as spoofing signals intelligence traffic. We should expect the net advantage to accrue to the side which enjoys a greater marginal benefit for the amount of resources invested. If AI notionally makes "finding" 10 times as cost-effective but renders "hiding" 100 times as cost-effective, we should anticipate hiding to become more effective in relative terms even though technology has radically improved the quality of finding in absolute terms. We argue below that theoretical analysis suggests that advances in AI could grant hiders a significant relative cost-effectiveness advantage under certain conditions. But if the two sides are mismatched in resources, then these marginal advantages might cancel out or even reverse. If "finding" is less cost-effective than "hiding" but the finder enjoys a massive advantage in total resources, "finding" can win out in practice even if in the abstract the technological environment advantages hiding.

[58] See, for example, David Barno and Nora Bensahel, "The Other Big Lessons That the U.S. Army Should Learn from Ukraine," *War on the Rocks*, June 27, 2022.

[59] Stew Magnuson, "Future Command's Rainey Eyes the Army of 2040," *National Defense Magazine*, October 6, 2023.

[60] Other analysts have expressed skepticism about a transparent battlefield for different reasons, focusing on the vulnerability of sensing systems to attack. Ivan Oelrich, Paul van Hooft, and Stephen Biddle, "Anti-satellite warfare, proliferated satellites, and the future of space-based military surveillance," *Journal of Strategic Studies*, Vol. 47, No. 6-7, November 9, 2024, pp. 916–39.

[61] Edward Geist, *Deterrence under Uncertainty: Artificial Intelligence and Nuclear Warfare*, Oxford University Press, 2023.

There are three caveats to this argument that suggest the hider-finder competition will remain especially contested, with neither side likely to gain an unequivocal edge.[62] First, the hider-finder competition interacts with the quality-quantity competition. If the finder can field an exponentially larger number of sensors than the defender can decoys, it makes it more likely that the finder can reduce the uncertainty that the deception campaign creates. Second, the amount of uncertainty the hider can create depends on the kind of information it is trying to conceal. It may be very difficult to prevent a finder from having a general sense of where large formations of the hider's forces are operating, but it may be much easier to prevent the finder from having exquisite intelligence on high value targets. This is especially true in the nuclear domain, where the finder's operational requirements are exceptionally high and its tolerance for uncertainty is likely at its lowest. Third, there is likely variation across different domains and physical environments. Wider and more cluttered search areas lend themselves to hiding, and smaller and emptier areas lend themselves to finding.

How AI Can Help Finders

The typical vision of how AI will help finders is through information fusion systems, which combine large amounts of data from various sources to build up an increasingly accurate and detailed picture of the state of the world. AI can help automate data-processing at a speed and scale that would be impractical for human teams of intelligence analysts.[63] For example, a new generation of small satellites already collects terabytes of imagery every day, and making sense of this much data within an operationally relevant timeframe is only possible with automation.[64] AI could also help finders field much larger numbers of sensors. For example, militaries could field "targeting meshes" with hundreds or thousands of drones equipped with various sensors.[65] Autonomy could strengthen the ability to operate large swarms of these drones without facing bottlenecks of skilled human pilots, preserve their ability to operate effectively in contested electromagnetic environments, and enable them to automatically recognize targets on their own. Taken together, AI and robotics could help militaries field larger and more capable sensing networks and analyze the data that flows through those networks more quickly and effectively.

From the finder's perspective, these advances in AI-enabled sensing can make kill chains faster and more lethal. Figure 3.1 illustrates the links of the so-called "kill chain." In 1996, Air Force Chief of Staff General Ronald Fogleman predicted that "in the first quarter of the twenty-

---

[62] Others have also used this "hider-finder" construct. See, for example, Krepinevich, 2015, pp. 65-66; and Antonio Calcara et al., "Why Drones Have Not Revolutionized War: The Enduring Hider-Finder Competition in Air Warfare," *International Security*, Vol. 46, No. 4, Spring 2022.

[63] Chang, 2021, p. 36.

[64] Saadia Pekkanen, Setsuko Aoki, and John Mittleman, "Small Satellites, Big Data," *International Security*, Vol. 47, No. 2, Fall 2022, pp. 177–216.

[65] Ochmanek and Hamilton, 2020.

first century, it will be possible to find, fix or track, and target anything that moves on the surface of the earth." With the addition of "engage" and "assess," these formed the links of the kill chain.[66] The usual intuition is that AI will strengthen kill chains by helping to connect the "find" and "track" link, as well as assisting in the "assess" link. Autonomous weapon systems and AI-enhanced decision-support systems could also help connect the "fix," "target," and "engage" links, such as by helping to manage battle networks that dynamically queue strike assets to engage time-sensitive targets once sensors locate them.

Figure 3.1 The Five Links of the Kill Chain

SOURCE: Geist, 2023, p. 85.

Some analysts view Ukraine as a preview of how AI can help make battlefields transparent. Because of the ubiquitous employment of small drones with sensors, coupled with advances in other methods like proliferated constellations of small commercial satellites, commentators have inferred from Ukraine that, "There is no more fog of war."[67] In this view, advances in robotics have generated a decisive data advantage for the finder, and AI will help solidify that advantage by quickly processing the data and rapidly acting on that targeting intelligence.

But this likely overstates the extent of the finder's current advantages and understates the countermeasures available to the hider. While AI will certainly help the finder, identifying a realistic upper bound on that improvement is important to have a realistic sense of how the measure-countermeasure race will evolve. Using AI to process and analyze information is relatively straightforward if the world is not too complicated and the data the sensing network collects are comprehensive and consistent. In practice, this is rare. The world is constantly changing, the information available is frequently misleading, often due to random noise and sometimes due to willful deceit. The flawed information that is available also only reflects a small subset of the overall state of the world. The AI information fusion system must consider all the different possible states of the world that might be consistent with the imperfect information that it has available and determine which of those states is most plausible. It has to make that determination with the knowledge that it certainly has incomplete information and that it may

[66] Geist, 2023, p. 85.

[67] "The added dangers of fighting in Ukraine when everything is visible: There is no more fog of war," *The Economist*, February 6, 2024.

also have inaccurate information. Therefore, information fusion is a form of reasoning under uncertainty.

Information fusion has inherent logical and computational constraints, which is where the problem starts for the finder. Reasoning under uncertainty is hard for both machines and people.[68] Because intelligent behavior often requires reasoning under uncertainty, AI researchers have spent decades attempting to find a way to make computers do it well. In the process, they developed an increasingly sophisticated theoretical understanding of reasoning under uncertainty. Several decades ago, this research revealed that a well-studied problem in theoretical computer science that is equivalent to information fusion—inference in Bayesian belief networks—is provably non-deterministic polynomial-time hard (NP-hard).[69] This means that no computationally tractable approximation algorithm can guarantee it will find a solution of the true value within a certain time interval. Additionally, it turns out that there is not one "right way" to represent and reason about uncertainty because there is more than one way to be uncertain. This means that one cannot overcome the obstacles to information fusion simply by employing a bigger computer, more sensors, or even some form of "superintelligent" AI.[70]

In sum, using AI to fuse together and analyze massive amounts of data can help the finder, but it requires reasoning under uncertainty. This process of reasoning under uncertainty presents structural challenges to how much finding can improve that simply developing more powerful computers and AI algorithms cannot solve. The key question for the hider-finder competition then becomes the conditions under which this task is most difficult for the finder and how a hider can best aggravate this challenge.

How AI Can Help Hiders

The key task for the hider is to complicate the finder's operational picture. The use of deception in war is so ancient that Sun Tzu deemed it paramount 2,500 years ago. We argue that hiders can offset advances in finding by using AI to develop "fog of war" machines. These AI-

---

[68] Edward Geist, "Why Reasoning Under Uncertainty Is Hard for Both Machines and People — and an Approach to Address the Problem," in Frank, Aaron B. and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, RAND Corporation, RR-A1275-1, 2022.

[69] Computer scientists sometimes use NP-hard as a shorthand for problems that are computationally intractable in the general case. NP is a complexity class containing those decision problems (problems with definitive "yes or no" answers) that can be computed in polynomial time on a nondeterministic Turing machine, a hypothetical type of computer that could evaluate multiple branches simultaneously. This class of decision problems is of particular interest to theoretical computer scientists because their answers can be checked "quickly" using a deterministic Turing machine, and, if P = NP, they may be the hardest decision problems that are computationally tractable using classical computers. NP-hard problems are those that can be reduced in polynomial time using a deterministic Turing machine to all problems in NP; therefore, a fast solution to any NP-hard problem is also a shortcut to solving every problem in NP. Some problems in NP are themselves also NP-hard; these problems are called NP-complete. Most NP-hard problems, however, are in complexity classes that include problems harder than those in NP.

[70] Geist, 2023, p. 128.

enabled tools would plan and implement military deception campaigns to maximize a finder's uncertainty, combining new technologies like robotic decoys as well as timeless military tactics such as operational-level feints and tactical-level concealment and dispersion.

The challenges of reasoning under uncertainty provide the hider with a potential structural advantage over the finder. All else equal, increasing uncertainty will be easier than resolving uncertainty. The hider's AI has an advantage in this competition because it needs to make it harder for the finder's AI to solve a complex computational task rather than solving one itself. Analogously, the finder needs to solve a puzzle, and the hider is making that harder by removing some puzzle pieces from the box, adding fake puzzle pieces at the same time, and trying to simultaneously change the picture so that solving the puzzle will still result in an outdated solution.[71] While the finder's sensors and information fusion system may improve, the adversary gets a vote in real-time on ways to degrade and pollute the operational the sensor picture that the finder's information fusion system sees.

Take the challenge for the attacker of closing kill chains as an example. Again, the attacker needs to accomplish multiple tasks to close the kill chain – find, fix, target, engage, and assess. The hider can use military deception to break the "find," "track," and "assess" links. For example, it can use camouflage and concealment to make it harder to find its forces; it can use mobility, decoys, and jamming the finder's sensors to make it harder to track its forces; and it can use emplaced decoys or spoofed signals to complicate battle damage assessments.[72] Military deception is a multi-level game where multiple players can employ different modes of deception on different scales. Hiders can degrade the finder's operational picture by creating imperfect information, inserting false information, or overloading the system with information.

An important tool of deception is to increase the adversary's uncertainty about the knowledge that it does have. While the ideal case for deception is that can prevent an enemy from gaining important information in the first place, such as preventing it from ever observing a task force of ships, that may become increasingly difficult as sensors become more ubiquitous. Instead, a fog of war machine offers particular value in increasing an adversary's uncertainty about whether the knowledge it does have is true and useful. If the hider has invested in large numbers of decoys, and if it coordinates the operations of real targets and decoys in ways that maximize uncertainty for the finder, then the challenge is figuring out what targets to attack and which are decoys. The finder's operational picture is likely to devolve into an extensive array of possible targets of varying probabilities of location and identification accuracy. The finder must grapple with uncertainties about the veracity and quality of targeting intelligence, scarcity of munitions, and cost-benefit trade-offs of exposing its own shooters.

---

[71] Geist, 2023, p. 126.

[72] Sherrill Lingel et al., *Leveraging Complexity in Great-Power Competition and Warfare: Volume I, An Initial Exploration of How Complex Adaptive Systems Thinking Can Frame Opportunities and Challenges*, RAND Corporation, RR-A589-1, 2021.

To better illustrate how the fog of war complex would work, consider the multiple ways to exploit deception for military effect. One approach is to create false information, such as fake forces to convince the enemy to misallocate its real resources. Allied efforts in Operation Fortitude during WWII exemplify this approach. The Allies created a fake Army Group to convince the Germans that the primary Allied landing would occur along the Pas-de-Calais rather than in Normandy.[73] Decoys such as inflatable planes and tanks and fake signals traffic accomplished this goal splendidly, misleading Hitler into misallocating his resources away from the defense of Normandy and helping ensure the success of the Allied landing in June 1944. But deception does not necessarily require using this kind of fakery.

Another approach to deception centers on the use of *real* forces to mislead the adversary about one's intentions. The Soviets favored this approach in WWII. Simultaneously with the Allied landing in Normandy, the Red Army embarked upon Operation Bagration against the German Army Group Centre. Soviet deception operations convinced the Germans that the next Soviet offensive would target Army Group North Ukraine rather than Army Group Centre.[74] This threat was very credible because there were actual Soviet armies posed to attack in Ukraine rather than Belarus.

Advances in AI could enhance both approaches to military deception individually and also create novel ways to combine them. The potential of AI to facilitate the creation of cheaper, more realistic decoys is obvious. To carry out the kind of deception used in Operation Fortitude, the United States created a "Ghost Army" staffed with actors and other creative types to undertake tasks such as faking plausible signals traffic.[75] Today, generative AI could replace thousands of human deception specialists with off-the-shelf software. Moreover, AI stands to make it possible to attempt this kind of deception with previously inconceivable speed and scale. Operation Fortitude required many months of planning and preparation, but in the future much of it might be actualized in weeks or days. With advanced robotics, the physical aspects of deception operations might be greatly accelerated as well. Decoy facilities could be designed, built, and "operated" by machines with minimal human input. As a consequence, their practicality and cost-effectiveness might be greatly enhanced relative to earlier wars.

Similar enhancements in cost-effectiveness might occur for deception operations utilizing actual forces. In quantity, unmanned systems can provide sufficient mass to devote a fraction to feints and other forms of deception while still retaining a large capability for the main effort. Inexpensive, autonomous unmanned platforms may also facilitate combinations of offense and deception that were not practical with past technologies. It would make very little sense to equip

---

[73] Robert Hesketh, *Fortitude: The D-Day Deception Campaign*. Overlook Press, 2000.

[74] Donald C. Daniel and Katherine L. Herbig. "Deception in Theory and Practice." In: *Strategic Military Deception*. Ed. by Donald C. Daniel and Katherine L. Herbig. Pergamon, 1982, pp. 355–367.

[75] Rick Beyer and Elizabeth Sayles, *The Ghost Army of World War II: How One Top-Secret Unit Deceived the Enemy with Inflatable Tanks, Sound Effects, and Other Audacious Fakery (Updated Edition)*, Chronicle Books, 2023.

an inflatable tank with a real gun. But a drone designed primarily for deception could be equipped with some real weapons (such as a warhead and the ability to be switched into an attack mode), forcing the adversary to take it seriously even if it knows that the drone is not what it pretends to be. Such combinations of deception and offense could permit novel concepts of operations.

Figure 3.2 Illustrative Fog of War Machine

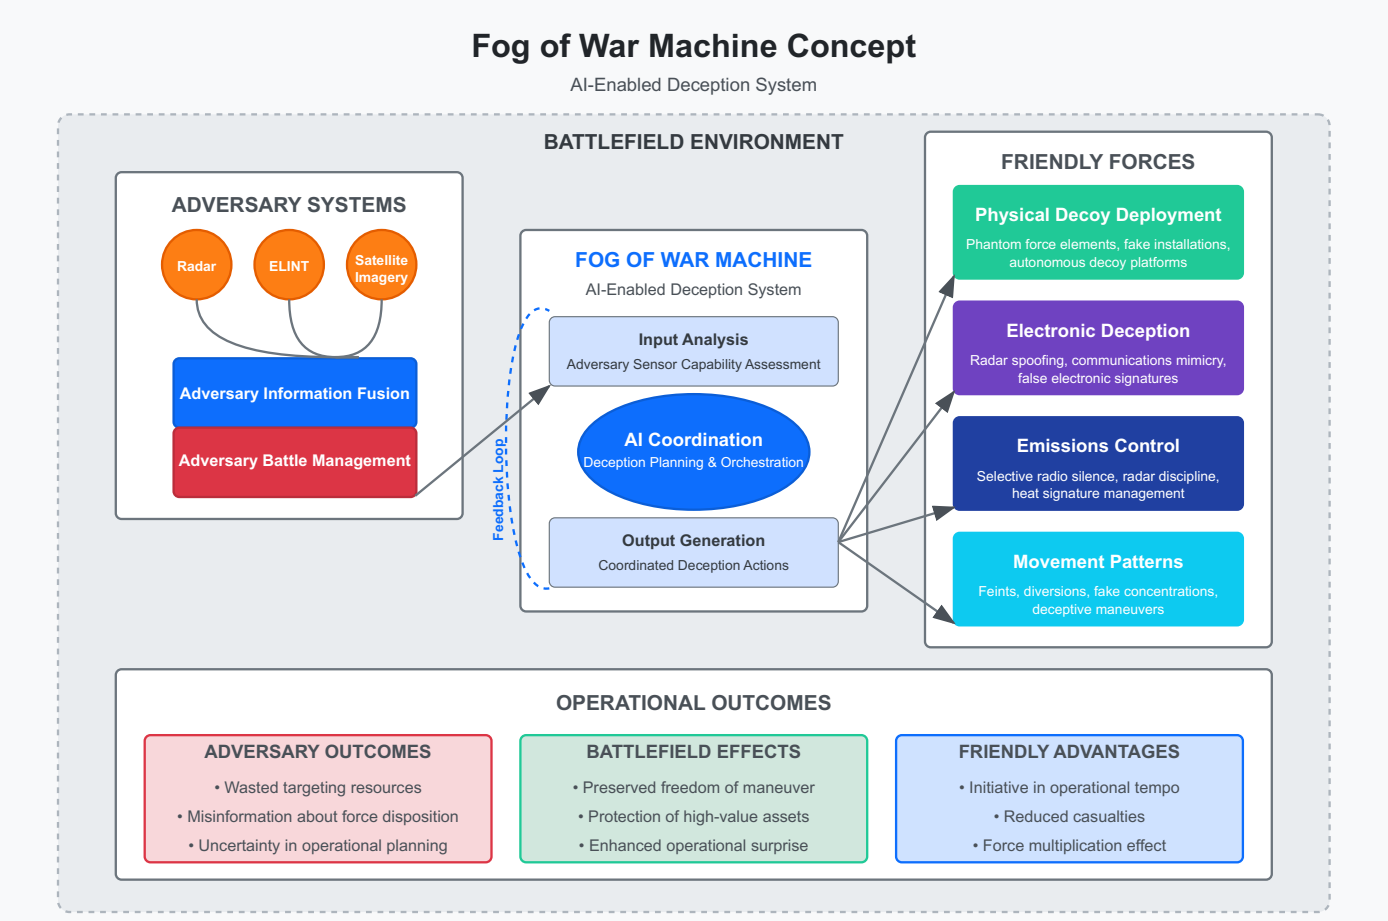

SOURCE: Ryan Bain produced this figure.
NOTE: In the figure, the arrow connecting the adversary's sensing systems to the friendly Fog of War Machine captures the idea that the Fog of War Machine would consider information about the adversary's sensing networks when planning deception operations. For example, if an imaging satellite is passing overhead, the Fog of War Machine might tip friendly forces to take cover to conceal their location or else queue a dazzler to suppress the satellite and mask the movement of friendly forces. In practice, friendly information about the adversary's information fusion and battle management systems will be imperfect and incomplete. Both sides might simultaneously be using Fog of War Machines to deceive the other, complicating efforts to optimize deception for an enemy's sensing network because that itself is subject to deceptive tactics to protect it.

Neither defenders nor attackers can be expected to enjoy a monopoly on these fog-of-war machines. Instead, we should anticipate that both sides in future conflicts might employ AI-enhanced military deception simultaneously. To sketch out a concrete example, imagine an

aggressor is attempting to conduct an amphibious landing against an island across a strait about 200 kilometers in width. The defending islanders aspire to defeat the invasion by sinking the aggressor's amphibious landing vessels as they cross the strait. To do this they plan to leverage a finite stockpile of anti-ship missiles to sink the ships. To reduce the effectiveness of these strikes, the attacker is incentivized to exploit military deception to incentivize the defenders to misdirect their fire in the hope that doing so will deplete the defender's supply of these munitions without sinking enough of the invasion force to forestall the landings.

A fog of war machine could exploit AI to plan and coordinate a dynamic deception campaign to thwart the defenders' target identification, tracking, and battle damage assessment capabilities. Well before the invasion, AI-assisted analysis could plan the design of the invasion force and its constituent vessels to facilitate the wartime deception effort. For example, the attacker might design invasion craft to resemble common civilian vessels, such as car ferries. In turn, this would enable the attacker to press civilian vessels into service as decoys if needed. The attacker could also design invasion craft to facilitate the development of decoys of varying degrees of plausibility. These decoys could be extremely elaborate with much of the complexity of a real ship, or they could consist of cheap fabric or inflatable superstructures floated upon buoyant Styrofoam foundations. They might even have no visible resemblance to an actual invasion craft, but rather consist of an unmanned surface vessel equipped with an electronic emitter able to feign the signature of an invasion craft based on the signals it transmits.

During the battle, the attacker would then seek to trick the defenders into squandering their supply of anti-ship munitions attempting to sink these decoys, while attempting as best they could to protect the actual invasion fleet. An extreme version of this strategy would mount a feint dash across the strait using solely decoys, perhaps with a few real high-value craft mixed in to give sufficient verisimilitude to an actual invasion. If the defenders took the bait, the attackers would make what appeared to be plausible mistakes—for example, intentionally leaking their "war plans" to the defenders with information that made the decoy fleet look like easy, tempting targets, particularly if they were targeted generously with anti-ship missiles. If the defenders began firing missiles at the decoy fleet, the attacker's fog of war machine could use AI to generate signals traffic indicating that the defenders' efforts to sink the invasion fleet was enjoying high success, and that if the defenders continued expending anti-ship munitions the invasion fleet would be sunk. The sacrifice of a few *actual* invasion craft and the capture of their crews within visibility of the defenders and world media could render this narrative plausible. Given the bias toward overestimating battlefield effectiveness in the moment and the illusion of imminent triumph, the defenders would face overwhelming temptation to fully commit their forces. But even if they understood they could be facing a trick, they might feel that they still needed to attempt to sink the invasion fleet even if there was only a minor chance it was real.

The defenders in this scenario, meanwhile, would have ample use for their own fog-of-war machines. The islanders would be launching anti-ship munitions from platforms they it would need to defend against attacks by the invader. They could deploy a fog-of-war machine intended

to mislead the invader about the location and status of these platforms. For example, if the invader was using anti-radiation and anti-aircraft missiles to attack air-based radar and missile-launching aircraft, they could coordinate their actual platforms with autonomous decoy UAVs to fool the invaders into wasting their missiles. They might employ AI-generated fictitious signals traffic to mislead the invader about the success of these attacks on the air platforms. They could undertake analogous maritime drone and signals spoofing efforts in the surface and underwater domains. A key goal of these deception efforts would probably be to cultivate uncertainty on the part of the invaders as to the total munitions inventory and salvo depth of the platforms in theater. This, in turn, could leave the invaders in the dark about the success of their effort to trick the defenders into wasting their anti-ship missiles.

Since both sides would presumably be anticipating the other's deception gambit, it would seem likely that both would only be partially successful. Both sides would attempt to draw in the other, only to have their opponent fail to take the bait. This could result in a lengthy "phony war" period of mutual deception and counter-deception. In short, tactical and operational deception would begat increasingly deep cycles of strategic deception. But as a crisis dragged on, one or both belligerents would plausibly come to feel increasing pressure for action, or could be convinced by chance factors that *this* time the adversary really was taking the anticipated action so that they had to counter it.

When Will AI Help Hiders or Finders More?

Fog of war machines are not a silver bullet for the hider-finder competition. Deception requires militaries to invest in decoys and to train in deceptive tactics, and then to employ those capabilities and tactics effectively in practice. This can involve tradeoffs. For example, emissions control (EMCON) is a valuable tactic to hide forces from a finder's sensing systems that can geolocate those emissions, but limiting how much they communicate or emit radar signals also imposes a tax on their operational effectiveness. This also applies to other deceptive tactics, like limiting movement and setting up camouflaged positions. These small inefficiencies could potentially generate cascading effects on their military utility. The investments in resources and training also create tradeoffs for force development, like investing scarce resources in decoys rather than more directly usable military capabilities.

There are three important variables that will shape the conditions under which hiders or finders gain an advantage in particular contexts. First, mass plays a key role. Mass can make deception more effective by creating so much noise that the finder cannot make sense of its environment, but sufficient mass on the finder's side might also neutralize some forms of deception. If the finder floods a designated area with enough sensors and aggressively destroys anything that has any chance of being a real target with masses of munitions, then it might end up destroying the targets of interest even if it is uncertain about which targets are real. But there are also significant constraints to a finder taking such a Carthaginian approach. The finder

typically has a finite number of responsive strike platforms and munitions, so using them recklessly would create capacity problems. The hider's employment of defensive interceptors and counter-battery fire might also further complicate the finder's ability to execute the kill chain at any time without putting its own forces at risk. Indeed, the hider-finder dilemma works both ways for a deceiver that has its own long-range strike complex and information fusion system.

Second, the amount of uncertainty the hider can create depends on the kind of information it is trying to conceal. A fog of war machine's greatest utility is creating uncertainty and confusion for the finder. In cases where it cannot actually "hide" the real target from the finder's sensors, it can create a shell game where the finder is uncertain which targets it finds are the real ones. This has the greatest potential to work for mobile systems, like road-mobile missile launchers and ships. It is unlikely to work well for hiding the locations of well-known fixed targets, like cities or ports. Additionally, it works best for creating uncertainty when the requirements for accurate and precise information are highest. If the hider wants to obscure the general position of an infantry brigade, that may be quite difficult. If it wants to prevent the finder from collecting targeting intelligence on a specific surface-to-air missile (SAM) battery, that is more feasible.

The nuclear balance is an area where the hider's task is comparatively easier than other areas. To enable a successful disarming strike, the finder needs to locate all or nearly all the hider's nuclear forces in a narrow timeframe and with a high degree of confidence.[76] Great powers with diverse nuclear force structures will have a range of different mobile systems that the finder needs to locate, including ballistic missile submarines (SSBNs), long-range bombers, and sometimes road-mobile transporter erector launchers (TELs). Unlike the conventional context, incremental attrition of these forces is of limited utility for the attacker. To pick notional numbers for illustrative purposes, if AI information fusion systems could allow the finder to go from locating 20 percent of the hider's conventional missile forces to 60 percent, that would be a dramatic and useful increase. In the nuclear context, it would still be a striking improvement but nevertheless fall far short of the requirements for a "war-winning" capability. Additionally, part of the goal of a fog of war machine would be to make the finder's human decision-makers unsure about where the percent they had located were, in fact, the real targets.

Third, there will likely be variation across different domains and physical environments. Wider and more cluttered search areas lend themselves to hiding, such as ground forces operating in urban or forested terrain. Urban terrain has both more physical site-blocking structures to hide from the finder's sensors and more activity that could create false positives where the finder may be unsure if it has found a civilian or military vehicle. Smaller and emptier areas lend themselves to finding. If forces are operating in narrow geographical areas, that reduces the search radius and how much uncertainty the hider can create. If they are operating in

[76] Absent radical breakthroughs in the effectiveness of missile defense, more than a handful of nuclear weapons surviving the attack may constitute a failure.

empty geographical areas, like ground forces in the desert or aircraft against the uncluttered background of the sky, that makes finding them easier.

Taken together, we can identify the conditions under which hiders or finders are most likely to gain an edge. The hider's job is easiest when it can better leverage mass numbers of decoys, when it is trying to prevent the finder from having high confidence about information that needs to be highly precise and accurate, and when it is operating in physical environments that are larger and more cluttered. The finder's job is easiest when it can better leverage mass numbers of sensors, when its goals require less precise information and enable greater tolerance for uncertainty, and when it is searching for targets in smaller and emptier physical environments.

Conclusion

Some analysts and policymakers believe that the future battlefield will be transparent. Central to this prediction is that AI provides a massive leap in information fusion capabilities, where it can collect and process data from a huge number of robotic sensors and quickly make sense of where targets are hiding. While advances in AI will strengthen the ability of militaries to find targets, it will also enable them to orchestrate deception schemes in much more sophisticated ways than they have in the past. Analysts too often overlook or dismiss the countermeasures available to the hider.

The analysis here points to a series of tradeoffs for both the hider and the finder. On the hider's side, practicing deceptive tactics like limiting movement and controlling emissions comes with an operational tax that will decrease their operational effectiveness. Investing in sophisticated physical and electronic decoys will also require allocating scarce resources that militaries might rather invest in platforms that are more directly combat capable. Because of advances in finding, the U.S. military will likely have to invest far more in paying this 'tax' to prepare for hiding than it has in the past. This will require a mindset shift from the past three decades when the U.S. military had overwhelming advantages of regional adversaries and insurgents where its main challenge was to find targets, not to avoid being found itself.

On the finder's side, there will be a different set of tradeoffs. At the tactical level, the finder will face tradeoffs about whether to strike targets when it is uncertain if they are actually the desired target. If the finder strikes all the potential targets that it locates, this induces two operational taxes. The first is that it reveals the finder's strike platforms to a hider's counterstrike capabilities. The second is that it could require a huge number of munitions, potentially creating capacity problems in terms of available strike platforms and the overall inventory of munitions.

The hypothesis that hiding will become more sophisticated and thereby keep the hider-finder competition contested in the future rests on the assumption that militaries take hiding seriously. Making deception effective is hard work, and finders could gain a serious edge if hiders fail to invest in the capabilities and tactics effectively. These techniques are made even more effective through schemes that leverage careful study of opponents to understand their baseline

operational approaches. All of this is difficult and costly. If hiders do not make sustained investments in sharpening their deception capabilities, the hypothesis may not hold.

Finally, the dynamics may look different at the operational and strategic levels of war. At the operational and especially strategic levels, militaries typically need to amass large arrays of forces for specific geographic or force-oriented objectives. Deception in this context also depends upon an adversary who is capable of being deceived. As Russia's preparations for invading Ukraine showed, it may be extremely difficult to hide large-scale war preparations, particularly in the ground domain.[77] But this hardly means deception becomes impossible. Deception depends as much on exploiting the finder's preconceived "conceptions" as it does the physical decoys, noise, or force disposition that the hider deploys.[78] While difficult, hiders have repeatedly demonstrated the ability to manipulate finder assumptions—regardless of the sophistication of the finder's information fusion systems—because deception retains a psychological component of ambiguity about intentions that is difficult to overcome.

---

[77] Jason Grey, "GEOINT Lessons Being Learned from the Russian-Ukrainian War," *USGIF*, February 22, 2024.

[78] The Hamas large-scale terror attack on October 7, 2022, illuminates the danger that even the most modern intelligence fusion enterprises face. Larry Hanauer and Michael P. Connell, "Political Priorities, Poor Intelligence Tradecraft, and the Suppression of Dissenting Views: Why Israel Failed to Warn of Hamas's October 7 Attack," Institute for Defense Analyses, December 2024. For classic texts on the crucial role of the target's assumptions and psychology, see Ephraim Kam, *Surprise Attack: The Victim's Perspective*, Cambridge, Mass.: Harvard Univ. Press, 2004; Uri Bar-Joseph, *The Watchman Fell Asleep: The Surprise of Yom Kippur and Its Sources*, Albany: State University of New York Press, 2005; Richard K. Betts, *Surprise Attack: Lessons for Defense Planning*, Washington, D.C: Brookings Institution, 1982; and Roberta Wohlstetter, *Pearl Harbor: Warning and Decision*, Stanford, California: Stanford University Press, 1962.

Chapter 4. Centralized vs. Decentralized C2

If AI can help humans make better and faster decisions, what are the implications for C2?[79] There are emerging debates about how AI could change C2 arrangements, with some envisioning a centralized "hive brain" and others foreseeing a host of decentralized AI agents.[80] The intuition is that AI could make automated decision-making more effective and reliable, reducing the need for humans to handle ambiguous and complex situations that overwhelm current-generation AI. China is reportedly interested in leveraging AI for centralized C2, although trust in AI systems remains a major political concern for the regime.[81] U.S. defense analysts have proposed using AI to enable more distributed and decentralized operations involving autonomous robotic platforms at a massive scale.[82]

We argue that the hybrid C2 approach of *mission command* will dominate. Mission command is an approach to C2 that involves centralized decision-making but empowers subordinate judgment at all levels about how to execute orders in ways that are best suited to the situation.[83] The hypothesis is that neither centralized nor decentralized C2 will gain a significant edge over the other and neither will become superior to mission command's hybrid approach. The central reason is that removing human intelligence as a constraint does not actually change the core reasons mission command is preferable. Mission command is currently dominant because it maximizes adaptability in the face of contingency and uncertainty and because no single actor has all the information necessary to make every decision. This enables a form of comparative advantage where higher-level commands can leverage their better informational picture of the overall operational situation and lower-level commands can leverage their better informational picture of their specific tactical situation. Increasing the reasoning ability of machines does not fundamentally change these dynamics, so it does not change mission command's relative

[79] On the limitations of current AI technologies for C2, see: Goldfarb and Lindsay, 2022; Walsh et al., 2021; and Cameron Hunter and Bleddyn Bowen, "We'll never have a model of an AI major-general: Artificial Intelligence, command decisions, and kitsch visions of war," *Journal of Strategic Studies*, Vol. 47, No. 1, 2024, pp. 116–46.

[80] See, for example, Sydney J. Freedberg, "Empowered Edge versus the Centralization Trap: Who Will Wield AI Better, the US or China?," *Breaking Defense*, February 13, 2024.

[81] Kevin Pollpeter and Amanda Kerrigan, *"The PLA and Intelligent Warfare: A Preliminary Analysis*", Center for Naval Analysis, October 1, 2021; Sam Bresnick, *"China's Military AI Roadblocks*", Center for Security and Emerging Technology, June 2024.

[82] Clark et al., 2020.

[83] This definition is an adaptation of U.S. Joint Staff and U.S. Army doctrine definitions for mission command. See U.S. Joint Staff, *JP 3-0, Joint Campaigns and Operations*, The Joint Staff, 2022, III-5; U.S. Army, *ADP 6-0, Mission Command: Command and Control of Army Forces*, U.S. Army, 2019, 1-3.

advantage. Rather than replacing mission command, AI is more likely to support and strengthen decision-making within this construct.[84]

How AI Can Help Both Centralized and Decentralized C2 Architectures

The prospect of advanced AI that can match or potentially even exceed human performance on cognitive tasks has clear appeal for C2 applications. Modern battlefields can be extremely complex, with large numbers of friendly and enemy platforms dispersed across wide geographic areas and in multiple domains. AI could help human commanders better understand their environments and make decisions about how to employ the forces available to them.[85]

Advances in AI could strengthen centralized approaches to C2. By centralized approaches, we mean theater-level commands exercising relatively tight and rapid direction control of forward forces. AI could strengthen the C2 process by providing decision-support to human commanders, such as helping them rapidly identify multiple candidate COAs and the requirements, risks, and potential enemy responses to those COAs. AI-enabled planning could be particularly useful for complex operations that require carefully sequenced effects across multiple domains. Much of AI's value for centralized C2 paradigms could also come from speeding up the implementation of decisions once higher-level commands make them. For example, translating strategic decisions into specific operational tasking can be a time-consuming process that AI could expedite.[86] AI could also facilitate faster dynamic adjustments in response to new orders or new information, such as updating operational logistics plans and Air Tasking Orders (ATOs) that currently face constraints in the ability of human staff to make adjustments relative to the tempo of operations.[87]

But advances in AI could also strengthen decentralized approaches to C2. By decentralized approaches, we mean that tactical command elements close to the forward edge of the battlefield have significant flexibility to make decisions about how they operate. AI might be particularly valuable for augmenting smaller command staffs with less planning experience. Forward units might benefit from AI-decision support tools to help them make sense of their environments and to fuse data on factors such as the disposition of enemy forces, terrain, and doctrine to help propose COAs.[88] Additionally, AI-enabled autonomy on the forward edge could be particularly important for C2 of unmanned systems operating in contested electromagnetic environments,

---

[84] Benjamin Jensen and Maj. Gen. Jake S. Kwon, "The U.S. Army, Artificial Intelligence, and Mission Command," *War on the Rocks*, March 10, 2025.

[85] As Chapter 3 explained, AI might also play a major role in increasing the prevalence and sophistication of deception, but all else equal humans will be better off trying to make sense of their environments if they have AI tools than if they do not.

[86] Clark, Patt, and Schramm, 2020.

[87] Walsh et al., 2021, pp. 4-6.

[88] Jensen and Kwon, 2025.

where the radiofrequency links required for humans to control them directly might be jammed. Similarly, autonomy will be particularly important for managing the operations of drone swarms where airspace deconfliction may be challenging and when there may be more platforms than there are human pilots available to operate them.

Mission command is a hybrid of centralized and decentralized C2 approaches. Central to mission command is the idea of "centralized decision-making, decentralized execution." This concept enablers higher-level commands to make decisions about strategy and key objectives, but it preserves the flexibility of forward units to determine the best way to pursue those objectives and to adapt as conditions change. As the next section outlines, U.S. military doctrine currently champions mission command, and we argue that it should leverage advances in AI to reinforce and strengthen this hybrid approach rather than changing course.

Why Mission Command Will Remain Dominant

A foundational assumption underlying our hypothesis is that militaries choose different C2 arrangements in response to tradeoffs. Drawing on Alan Docauer's research, we focus especially on the tradeoff between the proximity of forces to tactical engagements and their access to different kinds of information (Figure 4.1).[89] The key tradeoff is that theater-level commands have greater awareness of the operational environment, and forward tactical commands have greater awareness of the tactical environment. This information asymmetry means that neither centralized nor decentralized decision-makers have access to all the information they need about the best way to operate without relying on the other to some extent.

As distance from the forward edge of the battlefield increases, situational awareness generally decreases about tactical engagements and what is happening with the forces in contact with the enemy. Put simply, the farther you are from the front lines, the less you know what is happening on the front lines. By contrast, situational awareness generally increases about the overall operational-level environment because more attention and different data sources become available. That is, the closer you are to the front lines, the less you know and have time to think about what is happening across the entire theater. These types of information are not fungible because the knowledge that is necessary to make tactical decisions to prevail in a specific engagement is not the same knowledge that is necessary to make operational decisions about how best to conduct the campaign, which might include dozens of ongoing engagements with different needs and levels of priority.

[89] Alan Docauer, "Peeling the Onion: Why Centralized Control / Decentralized Execution Works," *Air & Space Power Journal*, Vol. 28, No. 2, 2014, pp. 24–44.

Figure 4.1. Situational Awareness vs. Proximity to Tactical Engagement (USAF)

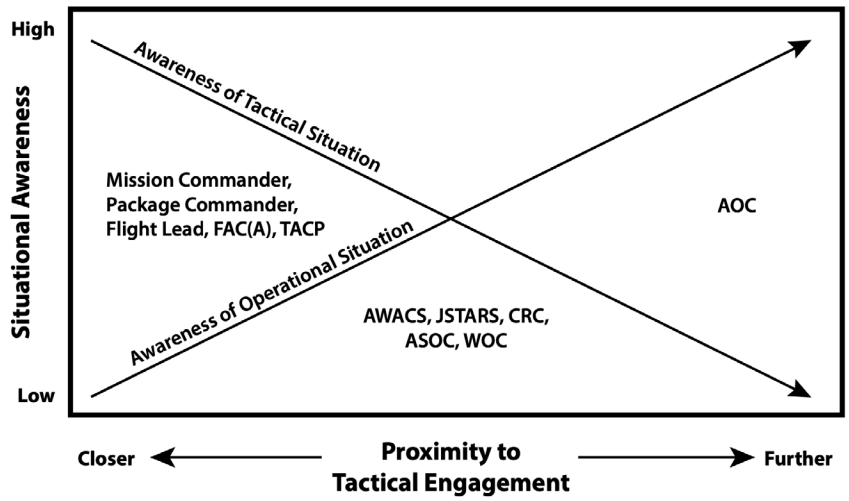

AOC - air and space operations center	FAC(A) - forward air controller (airborne)
ASOC - air support operations center	JSTARS - Joint Surveillance Target Attack Radar System
AWACS - Airborne Warning and Control System	TACP - tactical air control party
CRC - control and reporting center	WOC - wing operations center

SOURCE: Docauer, 2014, p. 27.

Mitigating these information asymmetries is structurally challenging. Long-range radio communications have given theater-level commands unprecedented reach since World War II, but human limitations in keeping up with the tempo and complexity of modern combat have not changed the incentive for some level of decentralized execution.[90] While this partly reflects asymmetries in the intellectual bandwidth of theater commands, which AI could help address, the constraints extend well beyond cognitive limitations. It is also hard for theater commands to have a clear picture of the tactical edge because the battlefield environment is dynamic and can change quickly, it is impractical for units to convey every piece of tactically-relevant information up the chain of command within a relevant timeframe, and warfare's inherent unpredictability makes it impractical to predict and orchestrate everything on the tactical edge in advance.[91] Theater-level commands may also lack tactic knowledge that is difficult to observe or communicate clearly, such as an intuitive sense of training proficiency or unit exhaustion. Additionally, establishing reliable communications links can be challenging against sophisticated adversaries because they can attack communications networks, such as jamming datalinks.

Current military C2 constructs have evolved partly in response to these tradeoffs. They consist of hierarchies of structure (e.g., theater, operational, and tactical echelons) and different

[90] Bailey, 2001; Anthony King, *Command: The Twenty-First-Century General*, Cambridge University Press, 2019.
[91] Walsh et al., 2021, p. 19.

functions (e.g., maneuver, fires, logistics, intelligence, protection) that require coordination up and down these echelons.[92] Theater commands decompose high-level orders into subordinate tasks that it assigns to a task force which in turn assigns them to a specific tactical unit.[93] The current C2 system operates with the expectation that there are significant constraints that will prevent theater-level and tactical commands from seamlessly sharing information with each other. Theater commands make decisions about force assignment and positioning orders at the operational level because they have a better view of the operational-level situation, while tactical commands make decisions about how to implement high-level orders because they have a better view of the many details required to understand their particular tactical situation.[94] The limiting factor driving this arrangement is not the cognitive capabilities of theater or tactical commands but their ability to access and act on different kinds of information.

Militaries still have a decision to make in how much authority they centralize versus delegate. The level of decentralization of C2 that a planner selects is often based on their needs in specific situations, including factors like the objective, commander's intent, threat, the rules of engagement (ROE), and anticipated contingency. Figure 4.2 visualizes how different considerations might shape how an advanced military like the United States might choose to centralize versus delegate certain tasks, but forces built around other principles and with different sensitivities might prioritize other factors in their decision-making matrix.

---

[92] Richard L. Daft, *Organization Theory and Design*, 10th ed (Mason, Ohio: South-Western Cengage Learning, 2010); Richard M. Burton, Børge Obel, and Gerardine DeSanctis, *Organizational Design: A Step-by-Step Approach*, 2nd ed (Cambridge University Press, 2011); Richard M. Burton and Børge Obel, "The Science of Organizational Design: Fit between Structure and Coordination," *Journal of Organizational Design* 7, no. 5 (2018).

[93] Take the task of crossing an enemy-defended river as an example to illustrate the value of nested echelon problem-solving. Planning for a river crossing under fire involves many technical details (river depth and speed, suitability of the crossing site, quantity and quality of bridging material) and intangible variables (subordinate commander experience and personalities, cross-unit familiarity and interoperability, and confidence in friendly counter-battery fire capabilities in a new situation). In this case, the division echelon is best positioned to plan a scheme of operations and make critical decisions due to its immediate access to dynamic physical data and its repository of tacit knowledge about subordinate commanders and force elements. Meanwhile, higher echelons at the corps level or theater level can focus on shaping the deep battlespace (by interdicting counter-attacking forces) and setting conditions (prepositioning supplies for reconstitution and maneuver forces for exploitation) so that the division's crossing will be successful. This interplay between centralized and decentralized C2 is what produces mission command – an ensemble of centralized plans, nested decision-making nodes, and decentralized execution.

[94] This gives theater commands primarily planning functions to design the overall campaign and force management functions to assure the right personnel and materials are positioned and sequenced to support the campaign plan.

Figure 4.2. Notional C2 Decentralization (USAF)

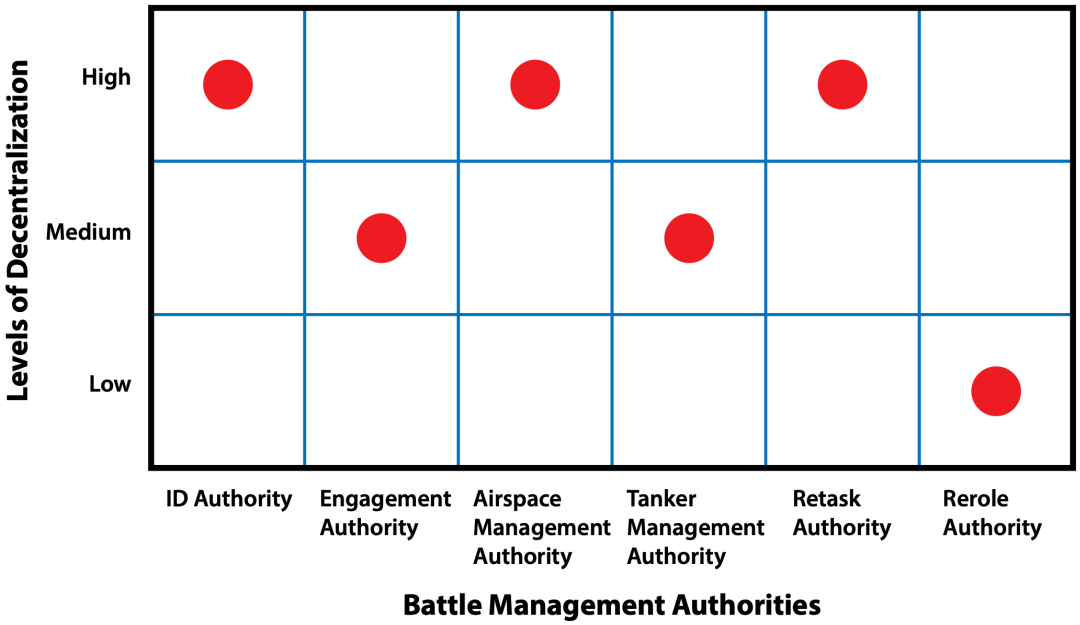

SOURCE: Docauer, 2014, p. 40.

Looking to the future, adding AI into this construct does not automatically change anything. Nor does anything really need to change. The sensors at the tactical edge do not necessarily have good views of broader areas. Communications links are all more vulnerable to jamming when operating close to the enemy, and the dangers of emitting at the tactical edge are significant because it can reveal the location of forces to the enemy. As militaries leverage AI in C2, the result might be the emergence of an ensemble of agents across various levels of action.[95] In this scenario, mission command even among an ensemble of AI agents remains a hybrid approach with centralized and decentralized aspects—for the same limitations of situational awareness, communications, and functional expertise that affect humans.

Because a key limiting factor for C2 is information rather than cognitive capacity, improving communications may yield greater dividends than improving intelligence. More robust communications networks can better share critical information to support decision-making. This includes for applications like swarms of robotic systems. Research on combat operations has consistently shown that big improvements result from adding small amounts of information between networked agents in swarms, something that is possible with even simple agents communicating with each other towards a goal and does not require a command element to

[95] K. Eric Drexler, "*Reframing Superintelligence: Comprehensive AI Services as General Intelligence*", Future of Humanity Institute, University of Oxford, 2019.

manage events at the edge.[96] In Figure 4.2 below, we assigned the attack element the task of attacking a cluster of targets by following a set of simple rules after arriving at the target area. This is a model of ordering a group to accomplish a mission without close direction that is consistent with mission-command arrangements. The results below come from a simulation of autonomous weapons with and without communications packages between platforms and allowing for tactical adaptation based on shared information.[97] The key result of the simulation is that information sharing enables adaptive swarm behaviors using smaller numbers of systems with simpler sensors and limited on-board decision-making. In Figure 4.3 the autonomous agents are adapting to increasing target location errors, and in Figure 4.4 the autonomous agents are adapting to attrition that triggers simple adaption noted as formation voids. In the case without communications, attackers approaching the targets are killed and not replaced, while in communication case the swarm adapts using a simple behavior that maintains swarm cohesion and coverage while dealing with a missing neighbor, all without need for centralized C2 to direct the attack and maintain tight control of the elements of the individual autonomous agents even in the face heavy attrition. The central insight is that intelligence is largely taking a back seat to communicating in improving combat outcomes.

Figure 4.3. Inter-flight Communications Increase a Swarm's Ability to Find and Kill Targets

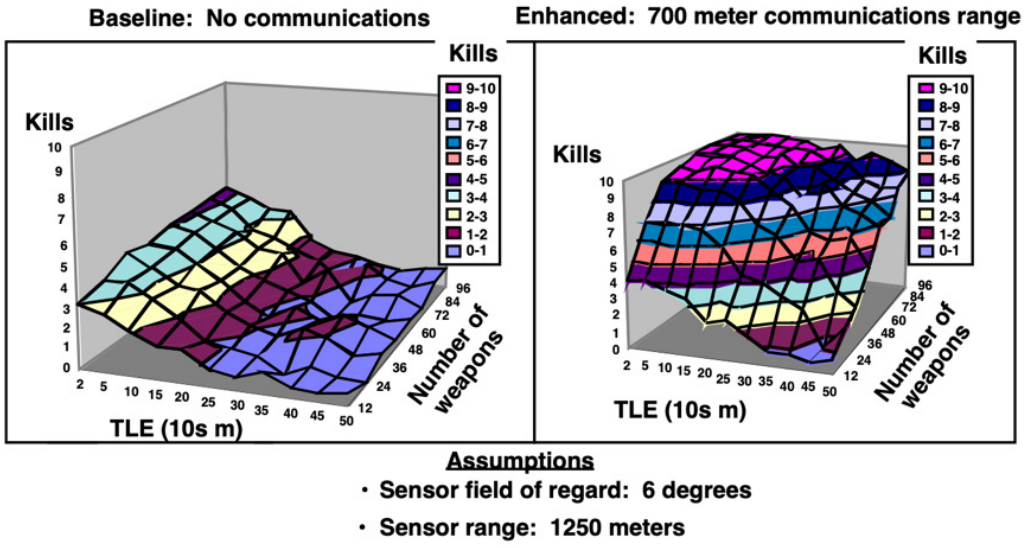

SOURCE: Frelinger, Kvitky, and Stanley, 1998, p. 41.

[96] David Frelinger, Joel Kvitky, and William Stanley, "Proliferated Autonomous Weapons: An Example of Cooperative Behavior," RAND Corporation, 1998. On a similar application to search and rescue, see: Ross D. Arnold, Hiroyuki Yamaguchi, and Toshiyuki Tanaka, "Search and Rescue with Autonomous Flying Robots through Behavior-Based Cooperative Intelligence," *Journal of International Humanitarian Action,* Vol. 3, No. 18, 2018.

[97] The model here assumes there is no error correction or synch overhead.

Figure 4.4. Inter-flight Communications Increase a Swarm's Robustness to Attrition

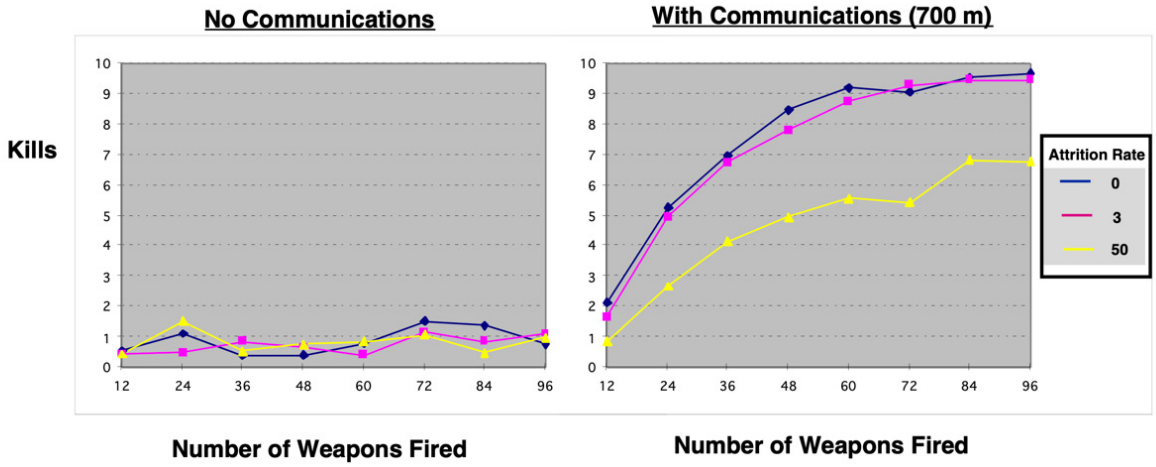

SOURCE: Frelinger, Kvitky, and Stanley, 1998, p. 43.

Conclusion

Centralized and decentralized C2 arrangements each offer value because different actors have access to different information. Theater-level commands have a better picture of the overall operational environment, and tactical commands have a better picture of what the forward edge of the battlefield looks like. Mission command works well because it combines the best of both approaches. The main limiting factor in making either centralized or decentralized C2 approaches inherently preferable is not intelligence but instead access to the right information, which is constrained by a range of factors including adversary attacks on communications links. Because better intelligence does not inherently solve this problem, the hybrid mission command approach will remain dominant.

Is AI truly transformative for C2? It will likely create important changes, but probably not to the degree people unfamiliar with military operations might imagine. Today, military forces operate with human general intelligences scattered from the edge of battle all the way to remote command and fusion centers, all with varying degrees of cybernetic assistance. AI might let people command autonomous platforms in ways that look more like how people command other people. However, the things that make tight control challenging and make exploiting information at the edge difficult apply to human and machine intelligences alike. In particular, the key is ensuring that communications pathways have sufficient bandwidth, latency, and protection when operating close to the adversary. AI might improve those systems in the design process or through fielding large mesh networks of autonomous communications nodes, but our physical reality creates some formidable obstacles to success, and a reactive adversary can also complicate things.

While AI does not change the dominance of mission command in relative terms, it does have the potential to improve both centralized and decentralized C2 architectures in absolute terms. AI's biggest benefits will likely be for helping build and operate the backbones of the C2 system itself, making communications networks more resilient, generating and evaluating COAs, and providing inferential estimates of possible enemy dispositions. This will all be particularly helpful for augmenting smaller command staffs that are constrained in the number of humans they have available to work these problems. There will also be a particularly important role for AI-enabled C2 in supporting autonomous systems in the field, where machines need to manage themselves or other machines if human operators are unavailable. But all these changes will likely happen within the mission-command construct.

Advances in AI could change the way that militaries currently implement mission command. The most significant attainable change to mission command may be AI-enhanced C2 enabling more *dynamic control* by helping manage tasks such as deconfliction, coordination, synchronization, and integration of complicated activities in a congested battlespace. To manage complexity, mission command typically relies on rigid mechanisms of predetermined control, such as pre-planned operational schemes, fixed unit coordination boundaries, or dedicated resource allocations. In the future, AI might help make real-time adjustments and reduce frictions that would have previously prevented *in situ* adaptation of complicated coordination schemes of control. This could support more dynamic and effective tactical responses in areas such as electronic warfare, airspace and fires deconfliction, and maneuver tactics.

But there will remain limits to how much AI-enhanced C2 will improve a military's overall combat effectiveness. Better decision-making and enhanced synchronization can help militaries make more efficient use of limited resources, but it cannot solve issues such as troops and equipment simply being unprepared for operations and unable to execute orders effectively, or time and space problems dictated by physical requirements such as maneuver and logistical sustainment, or inherent battlefield uncertainties about what the operational environment looks like and how best to act in response. AI will therefore likely strengthen C2, but it will not change the advantages that the mission command paradigm has, and improved C2 cannot overcome other physical and informational obstacles that limit military operations that are unrelated to the limits of human reasoning.

Chapter 5. Cyber Offense vs. Cyber Defense

How will AI affect the relative advantage between the cyber offense and the cyber defense? We argue that AI will improve both the cyber offense and defense in absolute terms, but it will likely create greater marginal benefits for cyber defenders relative to cyber attackers over the long term. The key reason is that AI could help address three key limiting factors that currently hold back the defense: speed, scale, and effectiveness. Defenders currently face resource constraints that limit their ability to discover vulnerabilities and defend networks relative to the size of the attack surface, and this gives attackers an advantage.[98] If AI can enable defenders to produce less error prone code in the first place, scan their software for vulnerabilities faster and more effectively, and develop and implement patches more rapidly and cheaply, this could disrupt the economics of cyber-security in a way that generally benefits the defender. Defenders could use AI to speed up and scale up these tasks in a way that erodes the advantages that attackers currently enjoy.[99]

We hypothesize that the cyber-defense will benefit more than the cyber-offense "in the limit" as the technology matures and diffuses. Any one computer network may remain vulnerable to attack even if the overall cost curve for the offense versus the defense in cyber-security changes. Additionally, the hypothesis is conditional on defenders taking advantage of the opportunities that AI provides by investing adequately in AI-enabled defenses. In the short term, if cyber-attackers leverage AI tools effectively but defenders do not, the offense's existing advantages could grow even larger.[100] But, in a long-term equilibrium where both sides leverage the technology with equal skill, we predict that advances in AI will benefit the defender relatively more by offsetting some of its current weaknesses. To be clear, the hypothesis is not that the defender will gain a structural advantage over the attacker in absolute terms, though that is possible under certain conditions. The hypothesis is that the defender will benefit relatively more

[98] For an overview of debates about the offense-defense balance in cyberspace, see: Rebecca Slayton, "What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment," *International Security*, Vol. 41, No. 3, Winter 2016/17, pp. 72–109.

[99] Other research has assessed the implications of AI for the offense-defense balance, but the analysis tends to focus on the short-term implications of AI whereas our hypothesis is about the long-term equilibrium of AI diffusion. Some analysts argue that neither the defender nor the attacker will gain an inherent advantage. Jenny Jun, "How Will AI Change Cyber Operations?," *War on the Rocks*, April 30, 2024; Andrew J. Lohn, *The Impact of AI on the Cyber Offense-Defense Balance and the Character of Cyber Conflict*, April 17, 2025, arXiv:2504.13371. Others argue that the balance of advantage depends on different factors. For example, after the release of new software, the attacker may have a short-term edge from faster vulnerability detection, but the defender may benefit over time because it is easier to find and patch all vulnerabilities given enough time. Ben Garfinkel and Allan Dafoe, "Artificial Intelligence, Foresight, and the Offense-Defense Balance," *War on the Rocks*, December 19, 2019.

[100] For an overview of early empirical examples of cyber-attackers using AI to support their operations, see: "Adversarial Misuse of Generative AI," *Google Threat Intelligence Group*, January 29, 2025.

from advances in AI to make its networks more resilient against cyber-attacks because AI can help address some of the weaknesses that currently hold defenders back relative to attackers.

To make this argument, we break down the cyber offense vs. defense building block into three component parts – vulnerability detection, maneuver and effects detection, and patching versus exploitation – and analyze how AI might reshape the dynamics of those competitions.[101]

How AI Can Help Both Attackers and Defenders Find Cyber Vulnerabilities

Gaining access to computer networks first requires the attacker to identify vulnerabilities in the defender's systems. Consider a basic model focused on software vulnerabilities.[102] Each system has N lines of code that contain V vulnerabilities, which are errors in the programming that an attacker can exploit. The number of vulnerabilities increases alongside the complexity of the system, since having more lines of code will tend to create more opportunities for errors, and some subset of those errors will result in vulnerabilities. Both the attacker and the defender are searching for vulnerabilities, which the attacker wants to exploit before the defender can find and patch. For the purposes of the model, assume that both sides are independently searching for vulnerabilities so that each side is essentially making a random draw from the set of V vulnerabilities that exist at any given time. The chance that the attacker or defender find a given vulnerability is therefore $1/V$ for each side. The chance that they both find the same vulnerability is $1/V^2$. When V is large, the probability that the defender finds the same vulnerability as the attacker is low. Closing the gap between what the attacker and defender will both find requires the defender to allocate vastly disproportionate resources as V grows larger. In sum, the attacker currently has a structural advantage in vulnerability detection.[103] Figure 5.1 below visualizes this advantage. The shaded grey region shows the difference between the probability that the attacker or the defender find a given vulnerability and the probability both find the same vulnerability.

AI could help the defender produce less error-prone code that has fewer vulnerabilities, thereby reducing the size of the attacker's asymmetric advantage. Following the model described above, if the number of vulnerabilities (V) decreases, then the chance that the defender finds the same vulnerability as the attacker increases ($1/V^2$). There are two opportunities to find and address vulnerabilities: making pre-deployment "fixes" on software before it is fielded or else making post-deployment "patches" to correct lingering errors.

---

[101] For a similar approach, see: Andrew Lohn and Krystal Jackson, "Will AI Make Cyber Swords or Shields?," *Center for Security and Emerging Technology*, August 2022.

[102] This model of software vulnerability is based on Ross Anderson, "Why Information Security Is Hard – An Economic Perspective," in *Seventeenth Annual Computer Security Applications Conference*, New Orleans, LA, 2001. For an alternative model, see: Ben Garfinkel and Allan Dafoe, "How does the offense-defense balance scale?," *Journal of Strategic Studies*, Vol. 42, No. 6, 2019.

[103] Anderson, 2001.

Figure 5.1. Baseline Cyber Offense vs. Cyber Defense

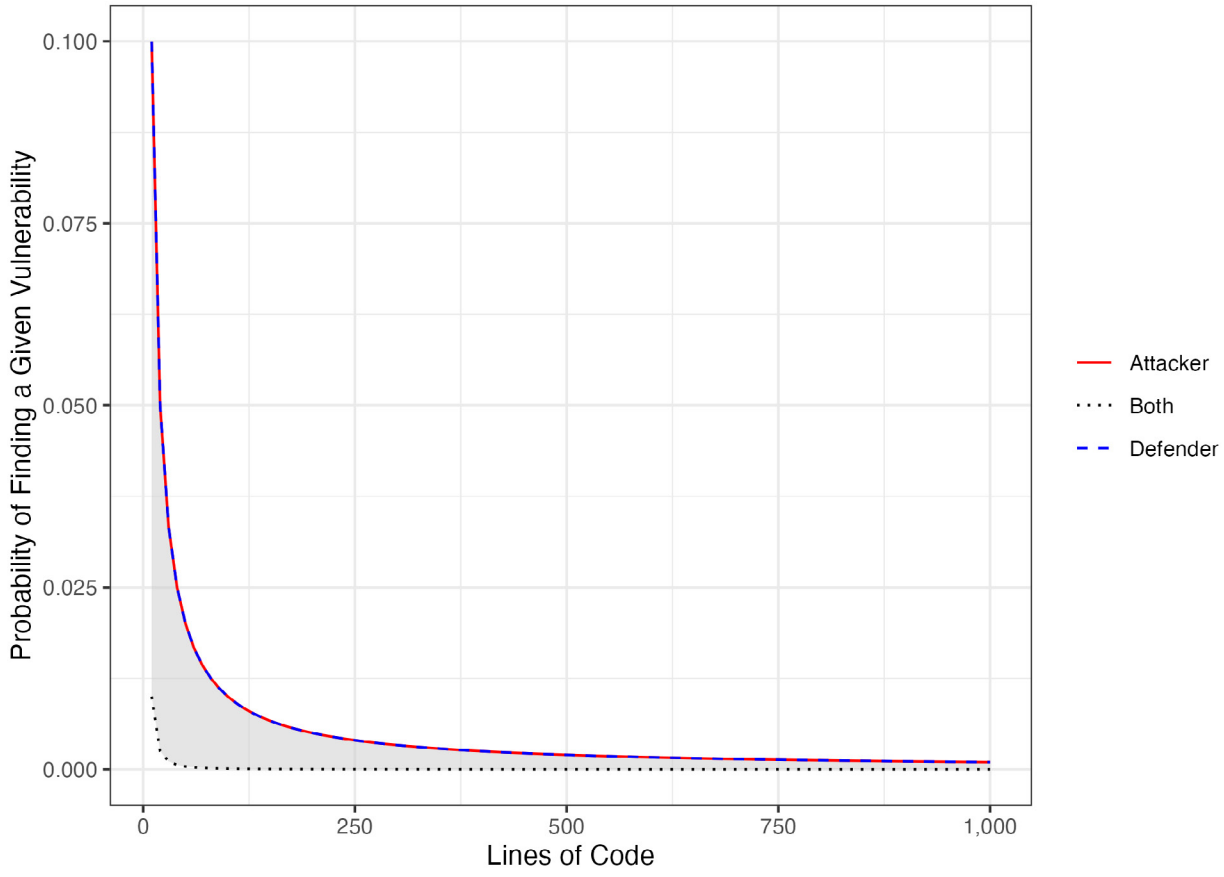

In the pre-deployment stage, AI could help the defender find and fix vulnerabilities before the attacker has the chance to start searching for them in the first place. This is a potential source of asymmetric advantage for the defender. If the defender can effectively use AI to produce more efficient code that is less prone to errors in the development stage, it could significantly reduce the size of the attack surface. The general rule of thumb today is that there will be 3 exploitable vulnerabilities per 10,000 lines of code (KLOC).[104] If the quality of AI tools for software development significantly improves in the future, *V* could start to decrease, and the attacker's current structural advantage would erode. Lowering the average rate of vulnerabilities to 3 per 100 KLOC or 1,000 KLOC, for example, would make the defender's job significantly easier. The attacker may always have some non-zero number of vulnerabilities that it could exploit, even if only because of problems with user errors or with the network's configuration, but the size of the potential attack surface could fall dramatically relative to what attackers enjoy today.

[104] Mark Sherman, "Using ChatGPT to Analyze Your Code? Not So Fast," *Carnegie Mellon University Software Engineering Institute Blog*, February 12, 2024.

Again, an important caveat to this argument is that the hypothesis is about shifts in relative advantage as AI matures and diffuses. In the short-term, the dynamics will be more complicated. AI is already playing an increasingly central role to software development. In 2023, 92 percent of software developers in the United States reported using AI coding tools.[105] In 2024, Google claimed that it had used AI to generate more than 25 percent of its new code.[106] But early research suggests that AI coding assistants and AI-generated code could potentially increase rather than decrease vulnerabilities in the short term, especially when implementation is poor.[107] Beyond AI coding tools simply generating code with mistakes, they are also creating new categories of problems: for example, AI sometimes hallucinates the names of software packages, so malicious actors have created fake packages with these names in anticipation of human programmers installing the AI-recommended packages without verifying their authenticity.[108] But if the technology continues to advance and humans learn to use it more safely and reliably, these problems may decrease after what could be a painful transition period. In the long run, it is even possible that AI agents will eventually write code and discover vulnerabilities at an expert level that outperforms most humans.[109] This could dramatically reduce pre-deployment vulnerabilities relative to their frequency today.

In the post-deployment stage, both the attacker and the defender can use AI to scan for vulnerabilities. The attacker still has an advantage in that it only needs to find one vulnerability to exploit, which is why it is so important for the defender to take advantage of opportunities to address vulnerabilities in the pre-deployment stage. Both attackers and defenders already use automated tools for vulnerability detection, but those automated tools still lag what expert humans can accomplish.[110] The attacker's ability to scale up its network reconnaissance could mean that a larger number of networks will come under threat compared to what attackers have the resources to target today. Attacking any one hardened network may therefore become much more difficult as well-resourced defenders gain better tools to reduce vulnerabilities before deployment and find them faster after deployment, but the overall amount of targetable attack surface across societies may increase, especially if there is uneven investment in AI-enabled cyber security tools.[111] This could increase the ability of militaries to use cyber-attacks against an

---

[105] Benj Edwards, "Google CEO says over 25% of new Google code is generated by AI," *Ars Technica*, October 30, 2024.

[106] Edwards, 2024.

[107] Jessica Ji et al., *Cybersecurity Risks of AI-Generated Code*, Center for Security and Emerging Technology, November 2024; Neil Perry et al., "Do Users Write More Insecure Code with AI Assistants?" *arXiv*, December 2023, arXiv:2211.03622v3.

[108] Sarah Gooding, "The Rise of Slopsquatting: How AI Hallucinations Are Fueling a New Class of Supply Chain Attacks," *Socket*, April 8, 2025.

[109] Lohn, 2025, pp. 11-12.

[110] Lohn and Jackson, 2022, pp. 8-9.

[111] Jun, 2024.

adversary's domestic critical infrastructure even as it decreases their ability to penetrate hardened computer networks dedicated to supporting specific military capabilities.

How AI Can Help Detect and Hide an Attacker's Maneuvers and Effects

Beyond finding vulnerabilities to enter the defender's networks, the attacker usually needs to maneuver within those networks once it has access. If the defender can detect the attacker's maneuvers or the effects that its operations are creating, this can alert the defender to the breach and prompt countermeasures. The attacker therefore wants to leverage stealthy cyber tactics to mask its operations, and the defender wants to detect the attacker's maneuvers and effects.

AI could help defenders better monitor their networks by improving their baseline understanding of their domain. Following the logic above that AI's largest benefits to the defender will be in the pre-attack phases, AI could strengthen the defender's ability to model and understand their networks better in ways that will bolster their ability to respond to attacks later. Currently, defenders often do not have a sufficiently accurate or complete view of their own networks, as artifacts like network diagrams can quickly become dated as networks evolve and normal behaviors like new hardware, software patching, and general maintenance accumulate over time. This may lessen the defender's understanding of their own network and create challenges in mounting an effective defense once targeted by an attacker. A focused attacker that conducts effective reconnaissance may even have a better view of the target's posture than the defender itself. As was the case for the resilience measures discussed above, AI tools may enable cyber defenders to much more quickly, cheaply, and effectively keep track of their own posture, increasing their preparedness to engage attackers once they detect an attacker on their networks.

Another argument often made in favor of the defense is that the improved baseline understanding of the defender's networks could increase its ability to detect attackers if it has a stronger sense of what activity on the network should normally look like. This can allow the defender to search for signs that the attacker has gained access by monitoring its networks and searching for anomalies, such as a certain part of the network activating at irregular times. Leveraging its ability to monitor its networks over long periods of time and to process increasingly large volumes of data, defenders could employ an AI-enhanced picture of baseline to identify deviations from a baseline, combining anomaly detection with known patterns of an adversary's operations. In theory, such a sophisticated AI-enabled monitoring and analytics of its networks at a speed and scale that is currently impractical or unaffordable with skilled human personnel may advantage the defender. In practice, however, there are also reasons to believe these approaches have fundamental limitations. Anomaly detection has long been a research topic in cybersecurity, finding mixed results.[112] AI-enabled anomaly detection in complex

[112] Joshua Steier, Erik Van Hegewald, Anthony Jacques, Gavin S. Hartnett, and Lance Menthe, *Understanding the Limits of Artificial Intelligence for Warfighters: Volume 2, Distributional Shift in Cybersecurity Datasets*, RAND Corporation, RR-A1722-2, 2024.

environments, especially those that may arise during crises or conflicts, has proven especially challenging as by its nature such environments do not confirm to past patterns of behavior, so everything may start to seem anomalous. False positives in such situations may be particularly problematic, diverting the defender's resources toward non-threats. Additionally, it is reasonable to expect that attackers will use the same technology to attempt to mask its maneuvers on a defender's network, learning from past successful attacks while also adjusting tactics so that it is more difficult for AI cyber-security tools to anticipate likely approaches. Finally, such approaches are necessarily reactive, and in a 'combat' environment where the attacker's objectives may call for swift action to target system availability (as opposed to cyberattacks focused on data confidentiality, such as the espionage attacks largely seen in the current environment) there are good reasons to believe that a detect-and-react approach may be insufficient to prevent many military cyber harms.

As a result, the relative advantage that AI offers the attacker versus defender in this sub-competition is less clear than in vulnerability detection, and it will likely remain largely a cat-and-mouse game. There is some reason to believe it might benefit the defender, especially in the longer term – but only if the defender does in fact leverage its home field advantage and uses AI to better design and understand its own networks during the pre-engagement phase. Similar to the hider-finder competition in Chapter 3, the more efficiently and securely the defender can use AI tools to reduce its attack surface and understand its operational environment, the harder and more complicated it will be for the attacker to stealthily operate on the target's networks for long periods, even if doing do so remains possible.

How AI Can Help Patch and Exploit Cyber Vulnerabilities

Finally, the attacker still needs to exploit vulnerabilities after it finds them, and the defender still needs to design and deploy patches to fix vulnerabilities. The delay between detecting a vulnerability and patching it gives the attacker a potential window of opportunity to achieve its desired objectives. How quickly the attacker can do so depends partly on if it needs to develop new code or if existing off-the-shelf tools work for exploiting a specific vulnerability.

AI could also benefit cyber defenders by speeding up their reaction time once they discover a vulnerability. On average, defenders are already slightly faster at writing patches than attackers are at developing exploits, but defenders face key limitations in the time it takes to deploy those patches.[113] Implementing the fix is by far the most expensive and time-consuming challenge for the defender. The defender has to not only develop the fix but also ensure that the update does not inadvertently cause other problems; for example, the cyber-security company Crowdstrike released an update in 2024 that interfered with other critical functions, resulting in millions of

[113] Lohn and Jackson, 2022, p. 11.

computer outages globally.[114] AI has the potential to execute this cycle at a much faster pace. AI could help not only write patches faster but also with implementation by speeding up testing of the patches with other critical software, generating the paperwork for change (e.g., requests, approvals, and records), and helping to maintain configuration plans. These are all tasks that make continuous defense so hard today, and AI has the potential to reduce these burdens that slow down the tempo and scale at which humans can operate. While the attacker might still retain some structural advantage in the probability of finding any one vulnerability before the defender, speeding up the development and deployment of patches could further offset that edge and enable the defender to mitigate the effects of vulnerabilities that do occur.

Conclusion

While commentators often focus on how AI will create new cyber-security risks and advantages for the offense, the defender may stand to gain the most from breakthroughs in AI. AI could help defenders find vulnerabilities faster, at a larger scale, and potentially more accurately than humans. It could enable faster response, analysis, and forensics of attacker movement, and operationalization of lessons learned across more systems, more quickly. AI could also automate the generation of a patch; the testing, verification, and validation of the system where the fix is required; and the necessary artifacts for re-accreditation and use. Taken together, these advances in speed, scale, and effectiveness could help cyber defenders address some of the challenges that currently hold the defense back.

In the long run, the cyber defense is only poised to benefit relatively more than the cyber offense from advances in AI if militaries and companies fully embrace the defensive potential of AI technologies and commit to adopting them. The attacker might still have an edge in the ability to attack any one network, but AI could help erode the size of the attacker's structural advantage.

AI could also create new kinds of vulnerabilities. In particular, militaries need to prepare for scenarios where cyber attackers try to seize control of the AI systems that the defenders build to protect their own networks. For example, if the defender creates an AI super user that has authorities to observe and change every element of their networks, then an attacker's ability to control that system could be catastrophic. That is, the AI itself becomes part of the attack surface. Ensuring the AI agent itself is safe and secure is essential to realizing the upsides of AI for the cyber defense without incurring tremendous risks. Especially in the early stages of fielding these AI agents for cyber security, this will likely require sacrificing some of its potential performance to create red zones around what it can do to certain systems and what it can do without humans in the loop.

[114] Lohn, 2025, p. 31.

Chapter 6. Does an AI First-Mover Advantage Matter?

The building block hypotheses examined in the preceding chapters assumed that both sides had developed and deployed advanced AI that could match or exceed human-level performance across a broad range of tasks (i.e., AGI). The main goal was to identify relative advantages *within* each building block assuming that peer competitors were facing off in a "mature" AI regime in which the technology had already diffused. But if one side has a first-mover advantage during a transition period, then the short-term equilibrium for the building block competitions may look different from the long-term equilibrium. Put differently, an AGI monopoly may offer the first mover the ability to defy the long-term advantages that we hypothesize and gain larger if fleeting advantages.

Historically, successful military innovations tend to spread, but there are often delays between when an innovation first appears on the battlefield, when it diffuses, and how it is adapted and adopted before a long-term equilibrium emerges.[115] This pattern suggests that there may also be a window during which one side has a first-mover advantage with AGI. There are three potential reasons that one side could gain a first-mover advantage. First, it could master the technology underlying AGI before its competitor manages to do so. Second, both sides may have the same technology, but the first mover is willing to accept greater risk in pursuing far-reaching military adoption of AI-enabled capabilities while the other holds back.[116] Third, both sides may want to adapt their forces, but the first mover may navigate the bureaucratic and acquisition process of replacing its legacy force structure much faster. That is, both sides may have access to the technology and share a preference to leverage it, but one side is better at executing that implementation.

The military implications of AI would differ during a transitional first-mover regime on the way toward a mature regime. Figure 6.1 outlines four possibilities depending on (1) whether "Blue" and "Red" maintain the status quo or adapt to AI and (2) when each side adapts. In the top right quadrant, Red embraces AI while Blue adheres to the status quo. During that transition period, Red could use its first-mover advantage in ways that might be impractical in a mature regime. For example, in the hider-finder competition, the finder may gain an exponential advantage if the hider has not yet developed a Fog of War machine and deployed decoys in mass.

[115] Michael Howard, *War in European History*, Oxford, UK: Oxford University Press, 1976.

[116] Horowitz, 2010.

Figure 6.1. First-Mover Advantage Typology

	Red Maintains Status Quo	Red Adapts to AGI
Blue Maintains Status Quo	Status Quo	"Transition" with Red First-Mover Advantage
Blue Adapts to AGI	"Transition" with Blue First-Mover Advantage	"Mature" AGI Regime

Having an AGI first-mover advantage does not automatically translate into decisive military advantages. From the perspective of the second mover, there are lower-tech countermeasures to offset some of the first mover's advantages even before it masters AGI. Those "lower-tech" countermeasures could even include ANI, which has already proven highly effective at performing specific tasks even if it sometimes has shortcomings in its reliability and adaptability. In the hider-finder competition, for example, the Fog of War Machine concept was originally envisioned as something that a second-tier power could use ANI to do to offset a leading military's AI-enabled sensor fusion.[117] Militaries have a long history of using deception to offset their weaknesses. For example, Serbian forces demonstrated tactics such as decoys and terrain masking in 1999 to partially offset U.S. and North Atlantic Treaty Organization (NATO) advantages in airborne and overhead sensing.[118] Much of the physical pieces required for a Fog of War Machine can be relatively basic, such as using shaped chicken wire to make a truck look like a mobile missile launcher or placing an electronic emitter on a small commercial drone. The second mover could develop an ANI that is tailor-made for orchestrating these military deception efforts, even if it lacks the ability to do much else. Red could also use existing jammers and missiles to attack the sensing and communications systems that support Blue's sensor fusion networks, polluting and degrading the data Blue's AGI receives in the first place.

The key implication is that *even if one side gains a first mover advantage, the second mover still has options to limit the extent of that advantage*, particularly if they are already using ANI. AGI does not inherently promise overmatch against opponents that lack frontier models. Moreover, if advantages accrue incrementally or with short measure/counter-measure cycles, the duration of that advantage may only be measured in months or, at most, a few years.[119] In fact,

[117] Geist, 2023.

[118] Benjamin Lambeth, *NATO's Air War for Kosovo: A Strategic and Operational Assessment,* RAND Corporation, MR-1365, January 1, 2001.

[119] Stillion and Clark, 2015.

the key limiting factor may be the ability, capacity, and resources that both sides have to produce mass quantities of key robotics systems rather than just how intelligent they can make them.

Whether the United States will have a first-mover advantage in AGI will depend as much on policy choices as technological breakthroughs. There are rational motives for militaries to be skeptical of calls for major change and innovation because such overhauls introduce uncertainty, and organizations dislike uncertainty.[120] We should therefore not assume that the U.S. military will automatically embrace AGI and gain a first-mover advantage in deploying the technology even if U.S. companies make key AGI breakthroughs before their international competitors.

The United States should not invest in new capabilities and concepts under the assumption that it will have a first-mover advantage, or that such a transition period would last a long time. Instead, it should focus on "no regrets" investments in approaches that will continue to work well even as the long-term equilibrium emerges.

[120] Barry Posen, "Military Doctrine and the Management of Uncertainty," *Journal of Strategic Studies*, Vol. 39, No. 2, 2016, pp. 159–73.

Chapter 7. Implications: A New AI-Enabled Way of War?

This working paper looked at four "building block" competitions to identify specific ways that advanced AI could influence military operations. The four hypotheses are that mass will have an advantage over quality, hiding will have an advantage over finding, mission command will have an advantage over more centralized or decentralized C2 arrangements, and the defense will have an advantage over the offense in cyber operations. Taken together, the central finding is that the characteristics of AI-enabled warfare will likely reward militaries that invest in capabilities and operational concepts that leverage mass, deception, mission command, and network resilience. Due to the deep uncertainty concerning how AI will continue to develop, these are tentative hypotheses.

Summary of Key Findings

First, AI could give quantity an advantage over quality. The core logic behind this hypothesis is that AI will enable militaries to field larger numbers of "good enough" unmanned systems at a much lower price point than manned systems. Because militaries will be able to generate more capability in total with the same resources, this will change the underlying cost-effectiveness of relying on quantity versus quality. Exquisite platforms and weapons will still have a role in future conflicts, but the relative advantages of mass will likely increase.

Second, AI could allow the hider to offset advances in finding. The core logic behind this hypothesis is that AI-enabled sensor fusion and large numbers of proliferated sensors could benefit the finder, but militaries can also leverage robotic decoys and AI Fog of War Machines to offset those improvements. All else equal, the hider has an advantage because reasoning under uncertainty is structurally challenging, and it is easier to generate uncertainty than to resolve uncertainty. But whether the hider or finder ultimately have an edge will depend particularly heavily on the context, such as each side's ability to leverage mass, the operational environment, and the kind of information that both sides are trying to uncover versus conceal. The result will likely be a contested measure-countermeasure race between hiders and finders.

Third, AI is unlikely to inherently benefit centralized versus decentralized approaches to C2, so the hybrid approach of mission command will remain dominant. The core logic behind this hypothesis is that the limiting factor in embracing more centralized or decentralized C2 models is access to the right information, not human intelligence. Forward and rear command elements often have access to different types of operational and tactical data that they need to support their decision-making. AI has a range of potential roles in C2, including to make communications networks more resilient and as a decision support tool, but it is more likely to reinforce the advantages of mission command as a C2 paradigm rather than to erode them.

Fourth, AI could benefit the cyber defense over the cyber offense. The core logic behind this hypothesis is that AI can help address some of the fundamental limitations in scale, speed, and effectiveness that currently hold the defense back. The task of defending very large attack surfaces against highly skilled attackers can place a large burden on the defenders that currently requires investing disproportionate resources in the defense to offset. The most important change that AI could generate for cyber-security would be enabling defenders to produce much less error-prone code in the first place, which would reduce the size of the exploitable attack surface and reduce the extent of the attacker's current asymmetric advantages. Cyber attackers will also benefit from using AI to support their operations, such as finding vulnerabilities and developing exploits, so attackers will retain the ability to penetrate networks at certain times and places, especially those that do not adequately invest in AI-enabled defenses. But AI will likely help reduce the magnitude of the attacker's advantages relative to the balance today.

Will these changes produce a revolution in military affairs? Our central predictions are about what characteristics AI-enabled warfare will reward, not about how to label those changes. Because AI is still maturing, there is a great deal of uncertainty about the technology's eventual capabilities and whether its effects will prove more evolutionary or revolutionary. If technologists are right that AI could eventually remove human intelligence as a constraint on completing cognitive tasks, and if militaries effectively adopt AI alongside supporting investments in other capabilities and technologies like robotics, that could have significant consequences for how militaries fight and win wars. As militaries consider this potential future, the analysis here suggests that they should explore capabilities and concepts that leverage mass, deception, mission command, and network resilience.

Implications for Conventional Warfare

If these hypotheses are right, what do they suggest about the future AI-enabled way of war? This section evaluates the implications of our hypotheses for long-term military competition between peers. We assume that neither side can afford to outspend the other dramatically, so cost-effectiveness is a key consideration.[121] The goal is to identify what kinds of military strategies, concepts, and capabilities are on the favorable side of the cost curve and what aspects of legacy approaches to war will become cost-ineffective and unsustainable. Militaries might still operate in ways that defy long-term trends, but they will end up paying a price for the privilege of doing so.[122] We also assume that both sides are operating in a mature AGI regime in which

[121] Stephen Biddle and Ivan Oelrich, "Future Warfare in the Western Pacific: Chinese Antiaccess/Area Denial, U.S. AirSea Battle, and Command of the Commons in East Asia," *International Security*, Vol. 41, No. 1, July 1, 2016, pp. 7–48.

[122] We do not seek to predict which countries will fight in which ways, because this depends on a host of political and strategic considerations beyond the capabilities that technology enables. See, for example, Caitlin Talmadge, *The Dictator's Army: Battlefield Effectiveness in Authoritarian Regimes*, Cornell University Press, 2015.

they need to make "bets" on strategies, concepts, and capabilities knowing that the competitor can employ its own AI-enabled countermeasures. Historical experience shows that pursuing approaches whose value depends on achieving and maintaining a long-term technological monopoly is likely to be a bad bet, so militaries should design their new way of war to work well even if the other side is just as technologically sophisticated.

If mass and deception will be central elements in the future of war, the United States will need to make significant changes to its existing force structure and operational concepts. Since the end of the Cold War, the U.S. military has prioritized quality over quantity and treated deception largely as an afterthought.[123] How might the United States redesign its force structure based on the principles of mass and deception?

Force planning should shift toward a much greater emphasis on mass while retaining a high-low mix. Embracing mass will require U.S. force planners to adopt a "Moneyball" approach to force planning.[124] Defense planners cannot rely on simply replacing one exquisite "star player" with a comparable next-generation exquisite system. Instead, the United States needs to look at how combinations of different systems can provide greater cost effectiveness than one system alone. This will require a mindset shift in terms of how the United States thinks about what attributes to prioritize in the procurement process. For lethality, a weapon's individual probability of kill may become less important than a salvo's probability of kill. For example, one Long Range Anti-ship Missile (LRASM) may be more capable than any single one-way attack drone, but a salvo of 10 AI-enabled drones may ultimately enable the United States to achieve a comparable battlefield effect at a fraction of the price.[125] For survivability, the United States will need to accept decreased performance and increased attrition to unmanned systems if it wants to lower costs and free up resources to build a larger force. This will require defense planners to resist the temptation to insist on performance and capabilities that mimic those of exquisite manned platforms, since this could effectively negative many of the potential cost savings of leveraging unmanned systems. Further research is necessary to determine what mix of attributes and performance requirements an AI-enabled force should have for different mission areas. However, the United States should not abandon exquisite capabilities entirely, especially in areas where it still has large advantages, such as undersea warfare. This will create tensions between continued investments in exquisite systems and new investments in mass.

Force planning must incorporate requirements for deception as a first-order requirement when developing new capabilities. Changes to how the United States approaches force planning and procurement will also be necessary to embrace the potential of AI-enabled deception. The current U.S. approach has often been to design its force structure around large, exquisite platforms and then only ask afterwards how it can hide them. This sequence makes

---

[123] Eaglen, 2022; Geist, 2023.

[124] Michael Lewis, *Moneyball: The Art of Winning an Unfair Game*, W. W. Norton & Company, 2004.

[125] Horowitz, 2024.

deception more difficult because it is harder to design cheap and effective decoys that mimic the size and signatures of exquisite platforms. Future force planning will need to consider deception as an important attribute when acquiring new capabilities. The spirit of lowering capability requirements to generate mass will also require the United States to avoid applying its traditional "gold-plating" mentality to decoys if AI can make larger numbers of good-enough decoys more effective than a few exquisite ones.

The shift to deception will require the United States to invest in Fog of War machines both as a metaphorical approach to warfighting and as a literal capability. The history of deception in war has been haphazard, with famous successes often depending on talented amateurs, but AI offers the opportunity to train a system that systematically optimizes how to create confusion and uncertainty persistently. Because deceptive tactics like camouflage and limiting movement or electronic emissions also impose a tax on U.S. forces, the Fog of War machine can also treat deception as a constrained optimization problem to balance competing priorities of operational effectiveness and survivability. The Fog of War machine will only be truly effective if it can orchestrate large numbers of physical and electronic decoys. This requires further analysis to figure out how to balance investments of scarce resources between new capabilities and decoys.

Militaries will need to counter adversary mass, not just field it themselves. In the long term, China could develop and field comparable or greater numbers of AI-enabled systems even if the United States started out with an advantage. In some cases, the United States might turn to symmetric responses, such as using drone swarms to defend against enemy drone swarms. In other cases, it will need to invest in asymmetric responses that put it on the right side of the cost curve. The war in Ukraine has demonstrated that electronic warfare is one potential option, with Ukraine reportedly losing 10,000 drones a month to Russian jamming. However, greater AI-enabled autonomy may reduce the vulnerability of drones to jamming. Another promising area for development is directed-energy weapons. Militaries have had long-standing hopes of employing effective directed-energy weapons, and these are now finally beginning to appear on mobile platforms on the battlefield.[126] Further breakthroughs in directed-energy weapons could conceivably provide a cost-effective way to counter large numbers of smaller, unhardened systems. While such breakthroughs are unlikely to stop the shift toward greater mass, they may limit the threat posed by small and cheap systems. AI is likely to play a large role in managing multi-layered air defense networks facing saturation attacks from an AI-enabled adversary.

Militaries will face a growing challenge in defending the lines of communication required to project power and sustain a mass-based force structure over long distances. Deploying an exponentially larger force structure would require developing much more extensive logistics capabilities to sustain that force. For example, if the United States builds

[126] See, for example, Venetia Gonzales, "Groundbreaking laser prototype systems delivered to 4-60th Air Defense Artillery Regiment," U.S. Army, September 21, 2023.

10,000 drones, it needs to be able to keep those drones supplied with fuel and ammunition.[127] The logistical ability to sustain force structures based on mass may be a limiting factor in how large military forces grow to be in practice, even with unmanned systems.

This will be especially difficult when militaries are projecting power far from home. Projecting power with a force structure based on mass will require not only transporting huge volumes of supplies to the right places over long distances but also protecting those lines of communication and logistical tail element from an enemy's long-range precision strike capabilities.[128] As U.S. competitors strengthen their ability to sense, target, and strike targets at long ranges, it creates new contested-logistics challenges.[129] Operating in mass will further complicate these challenges. Transportation connectors, such as container ships, oil tankers, trains, cargo aircraft, and long-haul trucks, may be particularly vulnerable as they traverse open terrain and predictable routes of travel. One option to address this challenge is to draw on concepts related to mass and deception. Using autonomous maritime platforms at scale, along with massed autonomous decoys, might prove effective enough to get on the right side of the cost-curve against the enemy's limited magazines of expensive, high-speed, long-range missiles.[130] Moreover, if these autonomous vessels can be supported with robotic over-the-shore connectors, then the vulnerability of a few deep-draft ports becomes less important.

Investing in a Defense Industrial Base that can build a force structure based on mass in the first place will be essential. The U.S. industrial base is currently struggling to build exquisite capabilities in limited numbers, such as producing only 1.2 attack submarines per year despite how important these capabilities are and the long-standing goal of fielding 2 a year.[131] Existing challenges with the defense industrial base that has evolved to create very complex weapon systems are both a constraint on how quickly the United States could build up its force structure quantitatively and a powerful incentive for the United States to invest in building smaller, cheaper platforms that the combined commercial and defense industrial base could more easily produce at scale. Another potential role for AI is to strengthen production techniques, such as enabling more effective robotic tools to generate higher-volume production and address

---

[127] Or else treat each as one-way attack drones or sensors, which would raise costs compared to reusable systems.

[128] Chris Smith, *"The Implications of Emerging Changes in Land Warfare for the Focused All-Domain Force*", Australian Strategic Policy Institute, September 2024, p. 19.

[129] Erik Sand, "Fighting to Supply the Fight—Assessing Approaches for Overcoming Contested Logistics," *Naval War College Review*, Vol. 77, No. 2, 2024.

[130] Smith, 2024.

[131] Alexander Grey, "The Submarine Workforce Crisis: Admitting Realities and Restructuring Long-Term Strategy," *War on the Rocks*, April 4, 2024.

bottlenecks stemming from the current need for a large, high-skilled human workforce.[132] But many key changes will require reforms to bureaucratic processes, not just technological fixes.[133]

The advantage may go to the side that *best executes the transition* to an AI-enabled force. In addition to determining the right mix of capabilities, a key area for future research is how to manage the transition between the force structure the United States has today and the one it needs to win wars in the future. One long-standing concern that has inhibited military innovation in the past is the fear that political authorities will call on militaries to fight when they are in the midst of incorporating a new innovation.[134] The concern is that being caught during the transition is the worst of both worlds, when the military is neither prepared to fight the way it knew worked in the past nor ready to fight in the way it believes will be important in the future. Managing this transition requires a plan for addressing the organizational cultural elements of adaptation, including thoughtful processes for how AI is fielded and integrated into units, monitoring how machines gain and maintain trust, and allowing for variations in how units use AI-enabled capabilities.[135] Successfully navigating the AI adoption process may be even more challenging and important than achieving the initial innovative breakthroughs with advanced AI, and it is far from clear that the United States will have an edge in the adoption competition.[136]

Implications for Strategic Stability

AI is unlikely to upend strategic stability between peer competitors. That is, the nuclear forces of both sides will likely retain the ability to withstand a first strike and retaliate.[137] All-out nuclear war is like Pokemon: "you gotta catch 'em all" to win. In conventional warfare, being able to destroy half of an enemy's arsenal of missiles, aircraft, or armored vehicles in short order would typically represent an enormous advantage. In the nuclear context, disarming 90 percent of the adversary's retaliatory forces might be insufficient absent significant improvements to missile defenses. In the hider-versus-finder competition, the extent of the hider's advantage will

---

[132] Patrick Tucker, "Could this breakthrough in robot manufacturing reshape warfare?," *Defense One*, April 14, 2025.

[133] These include changes to the planning, programming, budgeting, and execution (PPBE) process. Jeff Kojac et al., "PPBE Impact on Technology Transition: Findings and Recommendations," *White Paper Series*, Costello College of Business, Greg and Camille Baroni Center for Government Contracting, George Mason University, September 5, 2024. See also Commission on Planning, Programming, Budgeting, and Execution Reform, "Defense Resourcing for the Future: Final Report," Commission on PPBE Reform, March 6, 2024.

[134] Posen, 2016.

[135] Jonathan P. Wong et al., "One Team One Fight: Insights on Human Machine Integration for the U.S. Army," RAND Corporation, forthcoming.

[136] Radha Iyengar Plumb and Michael C. Horowitz, "What America Gets Wrong About the AI Race," *Foreign Affairs*, April 18, 2025.

[137] Kent, Glenn A. and David E. Thaler, *First-Strike Stability: A Methodology for Evaluating Strategic Forces*, RAND Corporation, R-3765-AF, January 1, 1989.

likely be greatest in the nuclear context, where the attacker must be confident that it can disarm all the defender's nuclear weapons for a counterforce attack to appear viable. Moreover, a counterforce strike requires near perfect target custody of the adversary's second-strike capabilities, not just for a brief moment, but likely for an extended period of time. The extreme operational requirements for a disarming strike leave little margin for error, and an AI fog of war machine is an ideal way to create enough mischief and mayhem that an attacker will feel less confident that it would reach that high bar for success if it were to attempt an attack. So long as one maintains sufficient doubt in the minds of potential attackers that their disarming strikes will operate flawlessly, then deterrence has a good chance of remaining robust.

There are other pathways through which AI could still generate instability and risks. First, AI could give the attacker advantages against much weaker adversaries. AI might not fundamentally change the threat to the survivability of great powers' nuclear forces, but it could create first-strike opportunities and temptations for great powers when facing regional nuclear adversaries with small arsenals. Second, militaries could misperceive the opportunities and constraints that AI creates. Attackers may be overconfident in the ability of their AI to see through the deception that the defender's AI orchestrates. Third, if the defender worries about the survivability of its nuclear forces, it could take destabilizing steps in a bid to reduce its vulnerability.[138] The most extreme version of this is giving an AI agent control to launch retaliatory nuclear strikes as part of an AI-enabled "dead hand." Given concerns about accidents and misuse, taking such a step would be a desperate measure and is unlikely. However, militaries could still use AI in other ways that appear safe in comparison but still increase the risk of accidents, such as deploying nuclear-armed autonomous systems on constant deterrent patrols.

Implications for Unconventional Warfare

Beyond traditional military operations, AI could open entirely new ways to assert control over a given piece of territory and its populace, particularly in situations where a more powerful AI-enabled state faced a weaker adversary instead of a great power peer. For example, could countries use AI to observe, harm, and influence others without physically occupying them? Such a "remote control" concept would involve one country using AI combined with robotics to extend its political control over key hubs of the target's society, such as its ports, airfields, industry, and critical infrastructure, as a means to exert coercive influence over the target.

A sensor grid and effector grid for remote population control. AI could provide the foundation of a sensing grid that collects and analyzes data from widely distributed sensors observing an area. States trying to exercise control over a target would likely find it is generally easier to observe the territory of a neighbor, since they enjoy the benefit of proximity. This can

[138] Michael C. Horowitz, Paul Scharre, and Alexander Velez-Green, "A Stable Nuclear Future? The Impact of Autonomous Systems and Artificial Intelligence," arXiv, December 2019.

include using a mix of sensors and intelligence types such as, imaging with electro-optical/infrared sensors and radar, signals intelligence to collect cellphone and other electromagnetic data, and cyber-attacks to collect personal data such as emails, credit card transactions, and social media data. Collecting this data over time and using AI to fuse and analyze it could allow the "control" state to build robust pattern-of-life analyses and create a population target data base.[139]

Beyond seeing, AI-enabled agents could also *act* to physically "control" a space. For example, small drones could monitor the flow of goods in and out of the target's ports. Drones with non-lethal capabilities could create mayhem by cutting cables, flattening tires, damaging cranes, and harassing dock workers. Beyond infrastructure, the virtual occupier could use reconnaissance drones and cyber surveillance to monitor the target's population, tracking what citizens are saying and where they are moving. This envisions an AI-enabled imitation of internal police state practices and asks how countries might translate them to a virtual occupation of external territories. The virtual occupier would have new tools to punish and incentivize certain behaviors – even down to the level of specific civilians – at a scale that reaches far beyond what is currently possible even with a physical occupation. AI could also help the occupier dynamically dial up and down these activities based on the target's compliance and to optimally plan its actions to have a maximal effect with a minimal footprint. The full range of these operations could extend from building a non-invasive surveillance net designed to create the feeling of a panopticon to a much more assertive "coercive quarantine" that aggressively employs air and naval drones to generate tailored costs on specific parts of the target's economy.

Additionally, countries could and presumably would extend these harassment operations to weakening the target's military. This could involve using the kinds of non-lethal effects described above to sabotage the target's military systems like aircraft at their bases and ships in port. Doing so could increase the target's vulnerability to conventional military coercion, supplementing the unconventional coercive approaches. Through a death by a thousand cuts approach, these subversive activities could extend effective political control over a target in peacetime and reduce the target's ability to defend itself if tensions boiled over into a conventional war.

Virtual occupation? If virtual occupiers could do this effectively, it may open the possibility that policymakers would be satisfied with a de facto virtual occupation without seizing territory by traditional physical means. Compared to making a conventional invasion easier or harder, AI in this vision would fundamentally change how countries think about what territorial control looks like at a first principles level.

But there are limitations and drawbacks to this kind of approach. First, the target could also use AI and drones to defend against the aggressor's activities. It could field drones and

---

[139] Yossi Sariel, *The Human-Machine Team: How to Create Synergy between Human & Artificial Intelligence That Will Revolutionize Our World,* eBookPro, 2021.

automated air defenses to protect airspace from foreign incursions, which might not protect every target all the time but would limit the scope and intensity of the hassling behavior the virtual occupier could orchestrate in practice. If the hypothesis is right that AI will benefit the cyber defender over the attacker, that would also constrain the virtual occupier's ability to use cyber-attacks and cyber-surveillance to tighten its grip over the target's society.

Second, the occupier would likely face a dilemma: the more effective its virtual occupation of the target proved, the more likely it would provoke retaliation and escalation if the target had the means to do so. The target could resort to symmetric attempts to use AI and drones, or it could turn to asymmetric retaliation, such as sanctions, political defiance, or threats to escalate to a broader conventional war.

This possibility illustrates that AI could shift the battlefield of international security competition into new and less familiar forms beyond its effects on altering the relative capabilities of attackers and defenders for traditional territorial conquest.[140] In this case, countries may gain a much more robust tool to harass and control each other, and may try to use those new tools as a substitute for traditional invasions and occupations.

Conclusion: Remaining Uncertainties

There is inherent uncertainty in making forecasts about technologies that are still emerging. However, there is already compelling evidence that AI will play important roles in the future of war, even if the most extreme visions of superintelligence and superweapons prove overblown. Ukraine has already demonstrated the value of applying current-generation AI to warfighting tasks like battlefield management systems and autonomy for low-cost drones.[141] The U.S. military has also made notable experimental breakthroughs, including a 2024 test in which an AI-controlled F-16 put up a "roughly even fight" against an F-16 with an expert human pilot.[142]

These demonstrations of AI's nascent capabilities in both active hostilities and peacetime experiments have two implications. First, AI's potential to have an important impact on military affairs does not depend on the assumption that the technology will experience exponential leaps in performance. Even if AI were only to improve linearly going forward, the question is still how big of an impact it will have rather than whether it will have a meaningful one. Second, adopting and integrating AI into U.S. military capabilities is a task that the United States needs to take seriously today, not defer until AGI is already at hand.

---

[140] Another key area to explore is the implications of AI for economic sanctions and non-military forms of cost imposition that states use to support their military strategies. On the complex dynamics that have historically shaped and constrained wartime trade policies, see: Mariya Grinberg, "Wartime Commercial Policy and Trade between Enemies," *International Security*, Vol. 46, No. 1, Summer 2021, pp. 9–52.

[141] Kateryna Bondar, *Understanding the Military AI Ecosystem of Ukraine*, Center for Strategic and International Studies, November 2024.

[142] Hadley, 2024.

The scenarios in which AI has the greatest potential to revolutionize military operations require more than just developing the technology. Militaries need to adapt their operational concepts and look for new ways to fight rather than merely using AI to make their existing approaches incrementally better. Unlocking AI's full potential also depends on other technologies, especially robotics. Without robots, AI is primarily an analytical and decision tool. When paired with robots such as unmanned aircraft or ships, it becomes much more versatile. Other technologies to monitor include advances in manufacturing techniques like 3-D printing to change the economics of producing robotic systems at scale, in energy storage to increase the performance and range of small platforms, and in directed-energy weapons to defend against swarms of robotic systems. AI might be at the heart of integrating all these advances together, but we should view it as part of a larger ecosystem and look for ways to cultivate and integrate supporting technologies into new operational concepts.

While there is a reasonably solid basis for projecting the direction of changes in relative advantages, there is greater uncertainty about their magnitude. There are compelling structural reasons to think that the cyber defense will get better and that hiding will become easier, but exactly how much more difficult cyber-attacks and reliably finding targets will become is far more difficult to predict. The extent of these advantages will also likely vary depending on the context, including different domains, types of missions, campaigns, and competitors.

Another remaining area of uncertainty is how quickly AI could lead to these advantages. There may be a striking difference in how militaries fight in the pre- and post-AI eras, but it may take decades for militaries to reach the new equilibrium. The pace of improvements could therefore make changes seem evolutionary even if the cumulative effects are eventually revolutionary. Even if one side developed advanced AGI first, it may still take considerable time to integrate AGI into its military systems at scale and to adapt its force structure to benefit from its military potential. Fully leveraging AI's potential to enhance robotic military capabilities may also require further advances in robotics and manufacturing to increase performance and reduce costs of those systems. There may be rapid progress in these fields, and AI could potentially accelerate those innovations, but the timeline of how long this would take remains unclear. Due to the remaining uncertainties, including over how quickly these hypothesized advantages could materialize, this issue deserves continued and detailed monitoring.

Abbreviations

AGI	artificial general intelligence
AI	artificial intelligence
ANI	artificial narrow intelligence
C2	command and control
CCA	collaborative combat aircraft
COA	course of action
ELINT	electronic intelligence
EMCON	emissions control
EO/IR	electro-optical/infrared
ISR	intelligence, surveillance, reconnaissance
JV 44	Jagdverband 44
LRASM	long-range anti-ship missile
NATO	North Atlantic Treaty Organization
NP-hard	non-deterministic polynomial-time hard
ROE	rules of engagement
SAM	surface-to-air missile
SSBN	ballistic missile submarine
SWaP	size, weight, and power
TEL	transporter erector launcher
UAV	unmanned aerial vehicle
UCAV	unmanned combat aerial vehicle
USAF	United States Air Force

References

Ackerman, Elliot, and James Stavridis, "Drone Swarms Are About to Change the Balance of Military Power," *Wall Street Journal*, March 14, 2024.

Allen, Gregory C., and Isaac Goldston, "Updating Augustine's Law: Fighter Aircraft Cost Growth in the Age of AI and Autonomy." Center for Strategic and International Studies, December 19, 2024.

Anderson, Ross, "Why Information Security Is Hard – An Economic Perspective," in *Seventeenth Annual Computer Security Applications Conference*, New Orleans, LA, 2001.

Arnold, Ross D., Hiroyuki Yamaguchi, and Toshiyuki Tanaka, "Search and Rescue with Autonomous Flying Robots through Behavior-Based Cooperative Intelligence," *Journal of International Humanitarian Action*, Vol. 3, No. 18, 2018.

Bailey, Jonathan B. A., "The First World War and the Birth of Modern Warfare," in *The Dynamics of Military Revolution, 1300-2050*, edited by Macgregor Knox and Williamson Murray, New York: Cambridge University Press, 2001, 132–53.

Bar-Joseph, Uri, *The Watchman Fell Asleep: The Surprise of Yom Kippur and Its Sources*, SUNY Series in Israeli Studies, Albany: State University of New York Press, 2005.

Barno, David, and Nora Bensahel, "The Other Big Lessons That the U.S. Army Should Learn from Ukraine," *War on the Rocks*, June 27, 2022.

Betts, Richard K., *Surprise Attack: Lessons for Defense Planning*, Washington, D.C: Brookings Institution, 1982.

Beyer, Rick, and Elizabeth Sayles, *The Ghost Army of World War II: How One Top-Secret Unit Deceived the Enemy with Inflatable Tanks, Sound Effects, and Other Audacious Fakery*, Chronicle Books, 2023.

Biddle, Stephen D., *Military Power: Explaining Victory and Defeat in Modern Battle*, Princeton, NJ: Princeton Univ. Press, 2004.

Biddle, Stephen, and Ivan Oelrich., "Future Warfare in the Western Pacific: Chinese Antiaccess/Area Denial, U.S. AirSea Battle, and Command of the Commons in East Asia," *International Security,* Vol. 41, No. 1, July 1, 2016, pp. 7–48.

Black, James, Matias Eken, Jacob Parakilas, Stewart Dee, Conlan Ellis, Kiran Suman-Chauhan, Ryan J. Bain, et al., *Strategic Competition in the Age of AI: Emerging Risks and Opportunities from Military Use of Artificial Intelligence*, RAND Corporation, RR-A3295-1, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA3295-1.html

Black, James, Rebecca Lucas, John Kennedy, Megan Hughes, and Harper Fine, *Command and Control in the Future, Concept Paper 1: Grappling with Complexity*, RAND Corporation, RR-A2476-1, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA2476-1.html

Bondar, Kateryna, *Understanding the Military AI Ecosystem of Ukraine*, Center for Strategic and International Studies, November 2024.

Bracken, Jerome, "Lanchester Models of the Ardennes Campaign," *Naval Research Logistics,* Vol. 42, Issue 4, 1995, pp. 559–77.

Bresnick, Sam, *China's Military AI Roadblocks*, Center for Security and Emerging Technology, June 2024.

Bristol, Henri, Enno de Boer, Dinu de Kroon, Rahul Shahani, and Federico Torti, "Generative AI in Manufacturing Lighthouses | McKinsey," McKinsey, April 2024.

Calcara, Antonio, Andrea Gilli, Mauro Gilli, Raffaele Marchetti, and Ivan Zaccagnini, "Why Drones Have Not Revolutionized War: The Enduring Hider-Finder Competition in Air Warfare," *International Security,* Vol. 46, No. 4, Spring, 2022, pp. 130–71.

Cares, Jeffrey R., Anthony Cowden, and Scott H. Swift, *Fighting the Fleet: Operational Art and Modern Fleet Combat*, Blue & Gold Professional Library, Annapolis, Maryland: Naval Institute Press, 2021.

Chang, Benjamin A., *Artificial Intelligence and the US-China Balance of Power*, dissertation, Massachusetts Institute of Technology, 2021.

Cohen, Rachel, "Perennial Pilot Paucity Puts Air Force in Precarious Position," *Air Force Times*, March 3, 2023.

Commission on Planning, Programming, Budgeting, and Execution Reform, "Defense Resourcing for the Future: Final Report," Commission on PPBE Reform, March 6, 2024.

Daniel, Donald C., and Katherine L. Herbig, eds. *Strategic Military Deception*, Pergamon, 1982

Daniels, Owen J., *The 'AI RMA': The Revolution Has Not Arrived (Yet),* The Andrew W. Marshall Foundation, October 2022.

Decker, Audrey, "Chinese satellites are breaking the US 'monopoly' on long-range targeting", *Defense One*, May 2, 2024.

Ding, Jeffrey, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, 1st ed., Princeton, N.J.: Princeton University Press, 2024.

Docauer, Alan, "Peeling the Onion: Why Centralized Control / Decentralized Execution Works," *Air & Space Power Journal,* Vol. 28, No. 2, 2014, pp. 24–44.

Drexler, K. Eric, *Reframing Superintelligence: Comprehensive AI Services as General Intelligence,* Future of Humanity Institute, University of Oxford, 2019.

Dwoskin, Elizabeth, "Israel Built an 'AI Factory' for War. It Unleashed It in Gaza," *Washington Post*, December 29, 2024.

Eaglen, Mackenzie, "The Bias for Capability over Capacity has Created a Brittle Force," *War on the Rocks*, November 17, 2022.

Engel, J. H., "A Verification of Lanchester's Law," *Journal of the Operations Research Society of America*, Vol 2, No. 2, May 1954, pp. 163–71.

Flournoy, Michele, "AI Is Already at War: How Artificial Intelligence Will Transform the Military," *Foreign Affairs,* Vol. 102, No. 6, 2023.

Forsyth, Robert, Jim Laurier, and Gareth Hector, *Me 262 vs P-51 Mustang: Europe 1944-45*, London: Bloomsbury Publishing, 2019.

Freedberg, Sydney J., "Empowered Edge versus the Centralization Trap: Who Will Wield AI Better, the US or China?," *Breaking Defense*, February 13, 2024.

Frelinger, David, Joel Kvitky, and William Stanley, "Proliferated Autonomous Weapons: An Example of Cooperative Behavior," RAND Corporation, 1998. As of April 22, 2025: https://www.rand.org/pubs/documented_briefings/DB239.html

Garfinkel, Ben, and Allan Dafoe, "Artificial Intelligence, Foresight, and the Offense-Defense Balance," *War on the Rocks*, December 19, 2019.

Garfinkel, Ben, and Allan Dafoe, "How does the offense-defense balance scale?," *Journal of Strategic Studies*, Vol. 42, No. 6, 2019.

Geist, Edward, *Deterrence under Uncertainty: Artificial Intelligence and Nuclear Warfare*, 1st ed., Oxford: Oxford University Press, 2023.

Geist, Edward, "Why Reasoning Under Uncertainty Is Hard for Both Machines and People — and an Approach to Address the Problem," in Frank, Aaron B. and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, RAND Corporation, RR-A1275-1, 2022. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA1275-1.html

Gerstein, Daniel M., and Erin N. Leidy, *Emerging Technology and Risk Analysis: Additive Manufacturing,* RAND Corporation, RR-A2984-1, February 15, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA2984-1.html

Goldfarb, Avi, and Jon R. Lindsay, "Prediction and Judgment: Why Artificial Intelligence Increases the Importance of Humans in War," *International Security*, Vol. 46, No. 3, February 25, 2022, pp. 7–50.

Gonzales, Venetia, "Groundbreaking laser prototype systems delivered to 4-60th Air Defense Artillery Regiment," U.S. Army, September 21, 2023.

Gray, Colin S., *Modern Strategy*, New York: Oxford University Press, 1999.

Grey, Alexander, "The Submarine Workforce Crisis: Admitting Realities and Restructuring Long-Term Strategy," *War on the Rocks*, April 4, 2024.

Grey, Jason, "GEOINT Lessons Being Learned from the Russian-Ukrainian War." *USGIF*, February 22, 2024. As of April 22, 2025: https://usgif.org/geoint-lessons-being-learned-from-the-russian-ukrainian-war/

Grieco, Kelly A., Hunter Slingbaum, and Jonathan M. Walker, *Cratering Effects: Chinese Missile Threats to US Air Bases in the Indo-Pacific,* Stimson Center, December 2024.

Grinberg, Mariya, "Wartime Commercial Policy and Trade between Enemies," *International Security*, Vol. 46, No. 1, Summer 2021, pp. 9–52.

Hadley, Greg, "In F-16 Dogfight, AI and Human Pilots Are 'Roughly an Even Fight,' Says Kendall," *Air and Space Forces*, May 8, 2024.

Hamilton, Thomas, and David A. Ochmanek, *Operating Low-Cost, Reusable Unmanned Aerial Vehicles in Contested Environments: Preliminary Evaluation of Operational Concepts*, RAND Corporation, RR-4407-AF, May 11, 2020. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RR4407.html

Hanauer, Larry, and Michael P. Connell, "Political Priorities, Poor Intelligence Tradecraft, and the Suppression of Dissenting Views: Why Israel Failed to Warn of Hamas's October 7 Attack," Institute for Defense Analyses, December 2024.

Hesketh, Robert, *Fortitude: The D-Day Deception Campaign*, Overlook Press, 2000.

Horowitz, Michael, "Artificial Intelligence, International Competition, and the Balance of Power - Texas National Security Review," *Texas National Security Review,* Vol. 1, No. 3, May 2018, pp. 36–57.

———, *The Diffusion of Military Power: Causes and Consequences for International Politics*, Princeton, N.J: Princeton University Press, 2010.

———, "Battles of Precise Mass: Technology Is Remaking War—and America Must Adapt," *Foreign Affairs,* Vol. 103, No. 6, 2024.

Horowitz, Michael, Paul Scharre, and Alexander Velez-Green, "A Stable Nuclear Future? The Impact of Autonomous Systems and Artificial Intelligence," arXiv, December 2019.

Howard, Michael, "The Forgotten Dimensions of Strategy," *Foreign Affairs,* Vol. 57, No. 5, Summer 1979.

———, *War in European History*, Updated ed. New York, NY: Oxford University Press, 2009.

Hughes, Wayne P., *Fleet Tactics: Theory and Practice*, Annapolis, MD: Naval Inst. Press, 1986.

Hunter, Cameron, and Bleddyn E. Bowen, "We'll Never Have a Model of an AI Major-General: Artificial Intelligence, Command Decisions, and Kitsch Visions of War," *Journal of Strategic Studies,* Vol. 47, No. 1, January 2, 2024, pp. 116–46.

Joint Staff, *Joint Publication 3-60, Joint Targeting*, Washington, DC: The Joint Staff, 2018.

Joshi, Shashank, "AI Will Transform the Character of War," *The Economist*, June 20, 2024.

Jun, Jenny, "How Will AI Change Cyber Operations?," *War on the Rocks*, April 30, 2024.

Kam, Ephraim, *Surprise Attack: The Victim's Perspective*, Cambridge, Mass.: Harvard Univ. Press, 2004.

Kent, Glenn A., David A. Ochmanek, Michael Spirtas, and Bruce R. Pirnie. "Thinking About America's Defense: An Analytical Memoir," RAND Corporation, September 29, 2008. As of April 22, 2025: https://www.rand.org/pubs/occasional_papers/OP223.html

Kent, Glenn A., and David E. Thaler, "First-Strike Stability: A Methodology for Evaluating Strategic Forces," RAND Corporation, R-3765-AF, January 1, 1989. As of April 22, 2025: https://www.rand.org/pubs/reports/R3765.html

Kessler, Sydney, Sam Wallace, Emily Yoder, and Jonathan P. Wong, "One Team One Fight: Ground Combat Narratives of Human Machine Integration," RAND Corporation, forthcoming.

King, Anthony, *Command: The Twenty-First-Century General*, 1st ed., Cambridge: Cambridge University Press, 2019.

Kojac, Jeff, Olivia Letts, Edward Hyatt, and Jerry McGinn, "PPBE Impact on Technology Transition: Findings and Recommendations," *White Paper Series*, Costello College of Business, Greg and Camille Baroni Center for Government Contracting, George Mason University, September 5, 2024.

Krepinevich, Andrew F., *Maritime Competition in a Mature Precision-Strike Regime*, Washington, D.C.: Center for Strategic and Budgetary Assessments, April 13, 2015.

———, *The Origins of Victory: How Disruptive Military Innovation Determines the Fates of Great Powers*, New Haven (Conn.): Yale University Press, 2023.

Lambeth, Benjamin S. "NATO's Air War for Kosovo: A Strategic and Operational Assessment." RAND Corporation, MR-1365, January 1, 2001. As of April 22, 2025: https://www.rand.org/pubs/monograph_reports/MR1365.html

Lepingwell, John W.R., "The Laws of Combat?: Lanchester Reexamined," *International Security,* Vol. 12, No. 1, Summer 1987, pp. 89–134.

Levenson, G. S., H. E. Boren, Dennis P. Tihansky, and Fred Timson, *Cost-Estimating Relationships for Aircraft Airframes*, RAND Corporation, R-761/1-ABR, January 1, 1972. As of April 22, 2025: https://www.rand.org/pubs/reports/R761z1.html

Light, Thomas, Fred Timson, and Obaid Younossi, *Parametric Flyaway Cost Estimating Relationships for Manned and Remotely Piloted Military Aircraft*, RAND Corporation, RR-1618/1-AF, September 26, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RR1618z1.html

Lingel, Sherrill, Matthew Sargent, Timothy R. Gulden, Tim McDonald, and Parousia Rockstroh, *Leveraging Complexity in Great-Power Competition and Warfare: Volume I, An Initial Exploration of How Complex Adaptive Systems Thinking Can Frame Opportunities and Challenges*, RAND Corporation, RR-A589-1, 2021. As of April 23, 2025: https://www.rand.org/pubs/research_reports/RRA589-1.html

Lohn, Andrew, and Krystal Jackson, "Will AI Make Cyber Swords or Shields?," Center for Security and Emerging Technology, August 2022.

Lynch, Christopher, Rachel Costello, Jacob L. Heim, Andrew Karode, Patrick Mills, Robert S. Tripp, and Alan J. Vick, *Operational Imperative: Investing Wisely to Bolster U.S. Air Bases Against Chinese and Russian Attacks*, RAND Corporation, PE-A1996-1, January 17, 2023. As of April 22, 2025: https://www.rand.org/pubs/perspectives/PEA1996-1.html

Magnuson, Stew, "Future Command's Rainey Eyes the Army of 2040," *National Defense Magazine*, October 6, 2023.

Menthe, Lance, Li Ang Zhang, Edward Geist, Joshua Steier, Aaron B. Frank, Erik Van Hegewald, Gary J. Briggs, Keller Scholl, Yusuf Ashpari, and Anthony Jacques, U*nderstanding the Limits of Artificial Intelligence for Warfighters: Volume 1, Summary,* RAND Corporation, RR-A1722-1, January 3, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA1722-1.html

Morgan, Forrest E., Benjamin Boudreaux, Andrew J. Lohn, Mark Ashby, Christian Curriden, Kelly Klima, and Derek Grossman, *Military Applications of Artificial Intelligence: Ethical Concerns in an Uncertain World*, RAND Corporation, RR-3139-1-AF, April 28, 2020. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RR3139-1.html

Oelrich, Ivan, Paul van Hooft, and Stephen Biddle, "Anti-Satellite Warfare, Proliferated Satellites, and the Future of Space-Based Military Surveillance," *Journal of Strategic Studies,* Vol. 47, No. 6–7, November 9, 2024, pp. 916–39.

Pekkanen, Saadia M., Setsuko Aoki, and John Mittleman, "Small Satellites, Big Data: Uncovering the Invisible in Maritime Security," *International Security,* Vol. 47, No. 2, Fall 2022, pp. 177–216.

Plichta, Marcel, and Ash Rossiter, "A One-Way Attack Drone Revolution? Affordable Mass Precision in Modern Conflict," *Journal of Strategic Studies,* Vol. 47, No. 6–7, November 9, 2024, pp. 1001–31.

Pollack, Kenneth, *Arabs at War: Military Effectiveness, 1948-1991*, University of Nebraska Press, 2002.

Pollpeter, Kevin, Amanda Kerrigan, and Andrew Ilachinski, *The PLA and Intelligent Warfare: A Preliminary Analysis*, Center for Naval Analysis, October 1, 2021.

Posen, Barry R., "Military Doctrine and the Management of Uncertainty," *Journal of Strategic Studies,* Vol. 39, No. 2, 2016, pp. 159–73.

Miranda Priebe, Alan J. Vick, Jacob L. Heim, and Meagan L. Smith, *Distributed Operations in a Contested Environment: Implications for USAF Force Presentation*, RAND Corporation, RR-2959-AF, 2019. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RR2959.html

Murray, Williamson, and Macgregor Knox. "Thinking about Revolutions in Warfare." In *The Dynamics of Military Revolution:1300-2050*, 1–14. Cambridge: Cambridge University Press, 2001.

Robinson, Eric, Daniel Egel, and George Bailey, *Machine Learning for Operational Decisionmaking in Competition and Conflict: A Demonstration Using the Conflict in Eastern Ukraine*, RAND Corporation, RRA815-1, 2023. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA815-1.html

Sariel, Yossi, *The Human-Machine Team: How to Create Synergy between Human & Artificial Intelligence That Will Revolutionize Our World*, eBookPro, 2021.

Scott Savitz and Amanda Perez, *Could the U.S. Navy Fleet of the Mid-21st Century Include Large Uncrewed Vehicles?,* RAND Corporation, PE-A2854-1, January 2025. As of April 22, 2025: https://www.rand.org/pubs/perspectives/PEA2854-1.html

Scharre, Paul, *Four Battlegrounds: Power in the Age of Artificial Intelligence*, First edition. New York: W.W. Norton & Company, 2023.

Schneider, Jacquelyn, and Julia Macdonald, "Looking Back to Look Forward: Autonomous Systems, Military Revolutions, and the Importance of Cost," *Journal of Strategic Studies,* Vol. 47, No. 2, 2024, pp. 162–84.

Sett, Gaurav, "How AI Can Automate AI Research and Development," RAND Corporation, October 24, 2024. As of April 22, 2025: https://www.rand.org/pubs/commentary/2024/10/how-ai-can-automate-ai-research-and-development.html

Slayton, Rebecca, "What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment," *International Security,* Vol. 41, No. 3, Winter 2016/17, pp. 72–109.

Smith, Chris, *The Implications of Emerging Changes in Land Warfare for the Focused All-Domain Force*, Australian Strategic Policy Institute, September 2024.

Stebbins, David, Richard S. Girven, Timothy Parker, Thomas Deen, Brandon de Bruhl, James Ryseff, Jessica Welburn Paige, et al. *Exploring Artificial Intelligence Use to Mitigate Potential Human Bias within U.S. Army Intelligence Preparation of the Battlefield Process*, RAND Corporation, RR-A2763-1, 2024. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA2763-1.html

Stillion, John, *Trends in Air-to-Air Combat: Implications for Future Air Superiority*, Center for Strategic and Budgetary Assessments, April 14, 2015.

Stillion, John and Bryan Clark, *What it Takes to Win: Succeeding in 21st Century Battle Network Competitions*, Center for Strategic and Budgetary Assessments, July 10, 2015.

Stillion, John, and David T. Orletsky. "Airbase Vulnerability to Conventional Cruise-Missile and Ballistic-Missile Attacks: Technology, Scenarios, and U.S. Air Force Responses." RAND Corporation, January 1999. As of April 22, 2025: https://www.rand.org/pubs/monograph_reports/MR1028.html

Sun Tzu, *The Art of War: Translated, with Introductions and Commentary by Ralph D. Sawyer,* Boulder: Westview Press, 1994.

Talmadge, Caitlin, *The Dictator's Army: Battlefield Effectiveness in Authoritarian Regimes*, Cornell Studies in Security Affairs. Ithaca London: Cornell University Press, 2015.

Tirpak, John A., "Kendall: In US-China 'Race for Technological Superiority,' AI May Be the Key," *Air & Space Forces*, October 29, 2024.

Turkes, Turker, *Fitting Lanchester and Other Equations to the Battle of Kursk Data*, thesis, Naval Postgraduate School, 2000.

Walsh, Matthew, Lance Menthe, Edward Geist, Eric Hastings, Joshua Kerrigan, Jasmin Leveille, Joshua Margolis, Nicholas Martin, and Brian P. Donnelly, *Exploring the Feasibility and Utility of Machine-Learning Assisted Command and Control: Volume 1, Findings and Recommendations*, RAND Corporation, RR-A263-1. As of April 22, 2025: https://www.rand.org/pubs/research_reports/RRA263-1.html

Washburn, Alan, *Lanchester Systems*, Monterey, CA: Naval Postgraduate School, March 1, 2000.

Watling, Jack, Oleksandr V. Danylyuk, and Nick Reynolds, *Preliminary Lessons from Ukraine's Offensive Operations, 2022–23*, Royal United Services Institute, July 18, 2024.

Wohlstetter, Roberta, *Pearl Harbor: Warning and Decision*, Stanford, California: Stanford University Press, 1962.

Wong, Jonathan P., Alexander C. Hou, Michael Miller, Katie A. Wilson, Emily Lathrop, Sydney Kessler, Sam Wallace, and Emily Yoder, "One Team One Fight: Insights on Human Machine Integration for the U.S. Army," RAND Corporation, forthcoming.

About the Authors

Zachary Burdette is an associate political scientist at RAND. He works primarily on research related to strategy, military analysis, and the intersection of economics and national security. Prior to joining RAND, he was a Research Fellow at the Harvard Kennedy School's Belfer Center for Science and International Affairs. He received a Ph.D. in political science from MIT, where he was a member of the Security Studies Program.

Dwight "Buzz" Phillips is a senior international / defense researcher at RAND. His research encompasses a wide range of areas including: defense strategy, operations, and readiness; technology, innovation and force development; force structure; mobilization and logistics; and civil-military relations. Prior to joining RAND, Phillips served in the U.S. Army in a variety of strategic planner roles for a combatant commander, service chief of staff, secretary of defense, chairman of the Joint Chiefs of Staff, and service secretary. Buzz received a Ph.D. in history from the University of Chicago.

Jacob Heim is a senior policy researcher at RAND. He specializes in strategic assessment and defense analysis with a background in international relations and mathematics. His research has focused on the military balance in the Western Pacific, defending air bases, USAF overseas force posture, and requirements for joint long-range strike. He holds an M.A. with honors from the Johns Hopkins University School of Advanced International Studies and a B.A. with distinction in mathematics from Amherst College.

Edward Geist is a senior policy researcher at RAND. His research interests include the former Soviet Union, nuclear weapons, emergency management, and artificial intelligence. Previously a MacArthur Nuclear Security fellow at Stanford University's Center for International Security and Cooperation and a Stanton Nuclear Security Fellow in RAND's Washington office, Edward received his Ph.D. in Russian history from the University of North Carolina in May 2013.

David R. Frelinger is a senior policy analyst at RAND. Since joining RAND in 1985, he has worked on a variety of research projects for the U.S. Army, U.S. Air Force, Department of Defense, Department of Homeland Security, the White House Office of Science and Technology Policy, and the National Geospatial Intelligence Agency. Frelinger's work has focused on developing appropriate responses to the development of AI, homeland security threats generated from the use of advanced and novel weapons, simulation and modeling, wargaming, serious games on non-national security issues, information technologies, and the interaction of commercial and governmental activities. David received an M.A. in political science from the University of California, Los Angeles.

Chad Heitzenrater is a senior information scientist at RAND. His work is centered on cyber warfare, signals intelligence, and mission assurance in the context of emerging warfighting

concepts. Heitzenrater's recent research examines the economics of information systems relative to secure software development, government acquisition, and cyber warfighting capacity. Chad earned his D.Phil. in computer science from the University of Oxford.

Karl P. Mueller is a senior political scientist at the RAND Corporation and a faculty member in the RAND School of Public Policy. He works on a wide range of national security subjects, including deterrence and coercion, airpower theory, U.S. grand strategy, alliances, security policies of small and middle powers, escalation and crisis stability, nuclear proliferation, space weaponization, and a variety of analytical and educational wargaming efforts. Karl received a Ph.D. in politics from Princeton University.